

Philosophy of Language and the Challenge to Scientific Realism

Christopher Norris

Contents

<i>Acknowledgements</i>	x
Introduction	1
1 Philosophy of language and the realism issue	6
2 The expert, the neophyte, and the X-ray tube: Hanson on ‘seeing aspects’	27
3 Philosophy of science as ‘history of the present’	50
4 Hanson on Wittgenstein, conceptual change, and quantum discontinuity	76
5 Saving appearances: the ‘linguistic turn’ and post-empiricist philosophy of science	98
6 Quantum mechanics and the limits of empiricism: recent challenges to the orthodox theory	120
7 Twin-Earth revisited: modal realism and causal explanation	143
<i>Notes</i>	175
<i>Index of names</i>	200

Introduction

This book investigates the sources and contexts of an approach to philosophy of science that emerged during the mid twentieth century and which achieved its most prominent articulation in Thomas Kuhn's classic *The Structure of Scientific Revolutions*.¹ My own approach is guided partly by the aim of reconstructing those sources and contexts, and partly by the need – as I see it – to treat them from the vantage point of recent debates on the topic of scientific realism. Those debates can often be highly bewildering for relative newcomers to the field since the term 'realism' is itself used in a great variety of senses, some quite technical or far removed from what it signifies in other, less specialist contexts. Likewise with the term 'anti-realism', which is sometimes applied in a general way to various opposed lines of argument, but sometimes to a certain logico-semantic (and metaphysical) thesis which denies the existence of unknowable or verification-transcendent truths.

Thus in the first chapter I offer a broad-based critical survey of various arguments for and against the realist (or objectivist) claim that truth-apt statements, in the physical sciences and elsewhere, have their truth-conditions fixed by the way things stand in reality rather than by our best current knowledge or even (at the limit) by our future best-possible means of proof or verification. After that, the focus switches to those earlier developments which may be seen to have set the philosophical agenda for a good deal of present-day discussion in these areas. Here the emphasis falls mainly on logical positivism/empiricism and on the various well-known problems with that doctrine which have continued to surface in the work of thinkers who often claim to have resolved them but whose solutions (I suggest) amount to little more than a shift of technical register. More constructively, I put the case for a full-fledged alternative approach to these issues in epistemology and philosophy of science, one that takes its cue from causal-realist theories based on a principle of inference to the best (most adequate or rational) explanation. Such theories in turn find support from certain recent developments in the field of modal logic, chief among them the Kripke–Putnam account of naming, necessity, and natural kinds. Thus my book takes the form of a running dialogue between realism and a range of anti-realist (or verificationist) positions, but one which – I should say straight off – comes out very strongly in favour of scientific realism. I trust that by the end readers will judge that this is

2 *Introduction*

not just a fixed prejudice on my part but an argument borne out by numerous case-studies in the history of science and also by reflection on the very nature and logic of scientific enquiry.

Hence my opening reference to Kuhn as the source of various post-1960 movements of thought whose chief common feature is their overt or covert opposition to the claims of scientific realism. Kuhn himself took a lead from Norwood Russell Hanson whose book *Patterns of Discovery* – published just three years earlier – had argued the paradigm-relativist case with constant reference to Wittgenstein. Thus Hanson treated the process of scientific theory-change as a matter of different observers perceiving different ‘aspects’ of the world in so far as their perceptions were theoretically informed and their theories underdetermined by the best empirical evidence to hand. Moreover, he pressed so hard on this analogy between paradigm-shifts and the *Gestalt*-based approach to issues of visual interpretation – as in Wittgenstein’s famous duck/rabbit example – that his argument lay open to a cultural-relativist reading which Hanson would surely have rejected had he lived to witness some of its more extreme present-day (e.g., ‘strong’-sociological) applications. Hence the unresolved tensions in his work, especially those that arise when Hanson brings this approach to bear on particular case-studies such as the debate between realists and strict empiricists with respect to what different observers ‘see’ when they witness (say) the phenomenon of anode fluorescence in an X-ray tube. I suggest that Hanson’s eager espousal of the Wittgenstein-inspired ‘linguistic turn’ leads him to adopt a framework-relativist stance which often gets sharply into conflict with his detailed reconstruction of historical episodes as well as his residual realist commitment when it comes to explaining just *how and why* such episodes occurred. Indeed this is among the most interesting features of Hanson’s work when viewed in the context of present-day arguments for and against scientific realism.

Thus my book adopts a critical-expository and sometimes diagnostic approach, drawing out these tensions and relating them to issues in more recent philosophical debate. I also discuss Hanson’s firm attachment to the orthodox (Copenhagen) theory of quantum mechanics, a theory which he takes – on the strength of phenomena such as wave-particle dualism – to require just such an aspect-relativist construal of previous episodes in the history of science. However, this position is upheld only through Hanson’s rejecting the claims of an alternative quantum theory (David Bohm’s ‘hidden variables’ interpretation) which perfectly matches all the empirical-predictive data while none the less affording a credible realist ontology and an adequate causal-explanatory framework. Through a series of further comparisons with (among others) W.V.O. Quine, Bas van Fraassen, and Richard Rorty, I emphasise the centrality of these issues to a range of developments in epistemology and philosophy of science since the demise of old-style logical empiricism. Hanson’s work can therefore be seen to prefigure both the emergence of various latter-day (post-Kuhnian) challenges to scientific realism and the problems that confront any cultural-relativist or strong-constructivist approach when applied to the long-run

‘context of justification’ as distinct from the immediate ‘context of discovery’. It also throws a sharply revealing light on the kinds of global anti-realist approach adopted by thinkers such as Michael Dummett and on the doctrine of ‘constructive empiricism’ – basically a more sophisticated variant of the positivist verification principle – defended with great ingenuity and resourcefulness by Bas van Fraassen. Thus my exposition of Hanson’s thought is interwoven with a series of critical commentaries on the sceptical or anti-realist turn in philosophy of science during the past three decades. My chief point throughout is that Hanson was possessed of a keen, knowledgeable, and deeply enquiring scientific mind which made him very much aware of problems – among them (intermittently) problems with his own approach – that often go unperceived by present-day thinkers of a similar persuasion. That they register more strongly in Hanson’s than in Kuhn’s treatment of kindred themes is one good reason for giving his work the degree of sustained exegetical attention that it has not received in recent years. Besides, as I suggest in chapters 1 and 4, this also might prompt a general reassessment of Wittgenstein’s currently powerful influence on the kinds of thinking that assimilate truth to what counts as such within some given community, practice, or cultural-linguistic ‘form of life’.

Hanson’s text can therefore be seen to have marked the most decisive turning-point in recent philosophy of science, the stage at which logical empiricism gave way – under pressures internal and external – to a range of alternative programmes which sought either to remedy its defects or to strike out in new directions. Some followed Quine in adopting a radically contextualist approach whereby scientific statements or theories are thought of as confronting the entire ‘tribunal of experience’ at any given time and we can always make adjustments at this or that point in the total web of belief, whether at the putative logical ‘core’ or the observational periphery. Others again took the Kuhnian path which led from Quine’s idea of ‘ontological relativity’ to that of the radical (world-transformative) effects of scientific paradigm-change, and thence – in some cases – to a strong-sociological approach which found no room for the standard distinction between ‘context of discovery’ and ‘context of justification’. However, there were those who saw the demise of old-style logical empiricism as a signal *not* for proceeding yet further in this sceptical direction but rather for asserting the claims of a realist approach to philosophy of science that would (in Wesley Salmon’s felicitous phrase) ‘put the “cause” back into “because”’. While Hanson’s book undoubtedly did much to encourage the former kinds of development, it also contains a good many passages (including some on the current situation in quantum-physical research) that can only be construed as appealing to the principles of causal realism and inference to the best explanation. It is on this account mainly – for its diagnostic value as a text which prefigures so many subsequent debates – that *Patterns of Discovery* merits such close attention. Above all it serves as a needful reminder of problems that are often concealed from view by sociologists of knowledge whose motivating interests – along with their fixed aversion to any form of scientific realism – incline them to ignore such strong counter-evidence from the record of scientific progress to date.

My own view – as will be evident by now – is that we cannot make sense of that record except on the causal-realist assumption which holds that successful scientific theories are those that are truth-tracking, or whose object-terms and predicates have picked out a constantly expanding range of genuine (real-world) objects and properties. Hanson explicitly rejects this view on Wittgensteinian, proto-Kuhnian, and orthodox quantum-theoretical grounds. However, it is one that may be challenged to strengthen its own philosophical defences through just the kind of sharply stated sceptical argument that Hanson brings to bear. Indeed this dialectic of realism and scepticism has been a prominent motivating factor at crucial stages in the history of science when a conflict arose between hitherto dominant and newly emergent paradigms. Thus the kernel of truth in Kuhn's account is the fact that major conceptual revolutions – like Copernican astronomy, relativity theory, and quantum mechanics – have often required a willingness to question what appeared the massive self-evidence of common-sense-realist thinking. Hence Einstein's early adoption of a positivist or Machian empiricist outlook that loosened the hold of such preconceptions and which thereby eased the passage to acceptance of a theory which would otherwise have faced much greater odds of deep-laid intuitive resistance. So likewise with previous revolutions in thought such as that which led some astronomers to doubt the plain evidence that the sun revolved around the earth or which allowed some nineteenth-century geometers to doubt whether Euclidean geometry was the only co-ordinate system that could possibly have any valid application. However, this context-specific case for the productive role of scepticism becomes much less convincing when extended – as it is by Kuhnians, cultural relativists, and strong sociologists of knowledge – into a full-scale programme for revising the history and philosophy of science in line with their favoured agenda. For it then goes beyond even Kuhn's idea that science alternates between relatively long-term periods of 'normal' problem-solving activity and relatively short-term, infrequent, and exceptional periods of 'revolutionary' ferment. Rather, so it is claimed, science should be characterised at best by a perpetual state of Kuhnian crisis where everything is open to creative 'redescription' and nothing stands firm against wholesale paradigm-change. Such is at any rate the argument of thinkers like Richard Rorty who count 'reality' a world well lost for the sake of multiplying novel descriptions and who regard any talk of scientific truth as a fallback to old, objectivist notions of a language-independent noumenal domain that somehow – impossibly – transcends our various currently favoured ways of describing it.

Einstein perhaps had a shrewd inkling of these later developments when he rejected Bohr's claims for the 'completeness' of orthodox quantum mechanics and abandoned his earlier positivist stance according to which there was no need – 'metaphysical' prejudice aside – to envisage a realm of causally operative real-world entities and forces beyond quantum appearances. Hanson not only accepted the orthodox theory but took it (along with certain lessons from Wittgenstein and cognitive psychology) to entail nothing less than a radical rethinking of the history of science right back to the Copernican revolution.

Whence Hanson's leading idea of the aspect-relative character of scientific knowledge and the impossibility of saying for sure that Tycho Brahe and Kepler might have seen 'the same thing' (i.e., the earth's axial rotation relative to the sun) from a hill-top location at dawn, given their respective allegiance to the old Ptolemaic and new Copernican worldviews. Such ideas have had a good run over the past half-century and provoked a great range of sophisticated arguments for and against. At present – to judge by the number and quality of recent books on the topic – there is a marked resurgence of interest in the case for causal realism and inference to the best explanation. Still, Hanson's text is a classic early statement of the paradigm-relativist approach and one that raises a strong challenge to both sides in this long-running debate.

I should perhaps mention that my opening and last two chapters make little or no explicit reference to Hanson, while chapters 2 to 5 keep his work constantly in view as a major point of reference. The reason for this is that I wanted to begin by bringing readers up to date with current disputes about scientific realism and to end by revisiting the same topics with the benefit of Hanson's insights, as well as the further insight to be had by reflecting on the various ambiguities, tensions, and unresolved problems in his thought. If this gives the book a somewhat sandwich-like structure I trust that it will offer adequate sustenance and not prove too indigestible.

1 Philosophy of language and the realism issue

I

Scientific realism takes various forms in various contexts of debate but mostly involves a commitment to the following propositions.¹ (1) There exists an objective reality which – *contra* idealists, ‘strong’ constructivists, and hard-line empiricists – is in no sense dependent on our thoughts, beliefs, descriptions, or theories concerning it. (2) These latter acquire their truth-value, that is, their status as truths or falsehoods, from the way things stand with respect to that objective reality and *not* from their happening to fall square with some currently favoured paradigm, conceptual scheme, or system of beliefs. (3) Among the vast (indeed limitless) range of truths about the world there are some that we know, some that we don’t but might yet find out, and some that may lie beyond the furthest reach of our perceptual, epistemic, or information-gathering powers. (4) Such truths obtain on every spatial and temporal scale, including (for instance) truths about the microstructural properties of matter, astrophysical objects and events, causal dispositions, laws of nature, historical facts (among them many unrecorded or unnoticed even at the time), prehistoric happenings right back to the origins of the universe, and so forth. Nevertheless (5), in so far as we can claim knowledge of them, that knowledge is acquired through our various procedures of observation, experiment, inductive reasoning, hypothesis-testing, or inference to the best (most adequate) causal explanation.² Such is the case for ‘convergent realism’, i.e., that if those procedures were not (for the most part) reliably truth-tracking then we could offer no account – short of a miracle or sheer cosmic coincidence – for the success of applied scientific knowledge in curing diseases, getting aircraft to fly, and a great many other achievements.³

However, it is still possible that we may be mistaken as concerns any given (no matter how well entrenched) item of scientific belief. This follows from the basic realist premise that truth is verification-transcendent and might therefore elude our present or even our future-best means of discovery. In which case, we are further obliged to accept that all knowledge is provisional (i.e., falsifiable) as a matter of epistemic warrant even though – or just because – the truth-value of our statements and theories is objectively fixed quite apart from whatever we currently think or believe. So there is no argument from the falsity – as we take it

– of so many past theories or truth-claims to the conclusion that truth cannot be conceived in objective or verification-transcendent terms.⁴ That idea is merely a result – so the realist will argue – of confusing ontological with epistemological issues, or questions such as: ‘what exists?’, ‘what are its properties, microstructural features, causal powers, etc.?’ with questions such as: ‘what are we presently in a position to assert (to the best of our knowledge) as regards those putative realia?’ Where the sceptic goes wrong is in thinking that the argument from past error leaves us with no rational option but to accept that scientific ‘truths’ are just those which happen to prevail within some given belief-community at some given stage of enquiry. What it shows, rather, is the need to acknowledge that now, as hitherto (and in the future), drastic revision might be called for should it transpire that our theories and truth-claims failed to match the best available evidence. Moreover – according to the realist – that evidence must itself stand under correction in so far as it might turn out to involve some perceptual distortion, anomalous result, or deficit in our range of empirical data. Still, this offers no support for the kind of sceptical meta-induction which would deny us any rational grounds for asserting that science makes progress or that our present-best theories are closer to the truth than those other, now discredited beliefs. For, again, such arguments ignore the fact that we *do* have strong corroborative evidence (if not ultimate proof) for the claim that science has enjoyed a large measure of predictive and causal-explanatory success and that causal realism is the only theory which doesn’t involve some miraculist appeal to cosmic coincidence or to our just happening to have hit on the right theories.

Anti-realism likewise comes in various strengths and philosophical forms but is also typified by a number of leading propositions. One is the denial that truth can be verification-transcendent, i.e., that we could ever have reason to maintain that some statement *x* must be objectively true or false quite apart from our possessing adequate epistemic warrant or our ability to produce evidence for or against it.⁵ Thus we are wrong to suppose that unproven mathematical theorems such as Goldbach’s Conjecture (that every even number is the sum of two primes) have an objective truth-value even though they cannot be conclusively checked by any means at our disposal and even though we can produce no formal proof that would obviate the need for such checking. Or again: we are wrong to think that certain unverifiable statements with regard to historical events – such as ‘George W. Bush uttered twenty-seven ungrammatical sentences *sotto voce* on the day of his presidential Inauguration’ – must be either true or false despite the unlikelihood (indeed the near impossibility) that such evidence should ever turn up. In cases like these, so the argument goes, we are unable to acquire or to manifest a grasp of their operative truth-conditions, and must therefore concede that they lack any kind of genuine assertoric warrant. In its full-strength version, as espoused by anti-realists like Michael Dummett, this entails the idea that any ‘gaps in our knowledge’ must be taken as corresponding to ‘gaps in reality’ since reality so far as we can possibly know it is coextensive with the range of well-formed statements for which we possess some proof-procedure or adequate means of ascertainment.⁶ Thus, for statements of the so-called ‘disputed class’ (like those instanced above),

we had best accept that the principle of bivalence fails, i.e., that they are neither true nor false, or at any rate – on Dummett’s more qualified version of the thesis – that we are not in a position to assert that they must have some objective truth-value albeit unbeknownst to us. This he takes to follow as a matter of logical necessity if one accepts Frege’s cardinal precept that ‘sense determines reference’ along with the principle that the meaning of a sentence is given by its truth-conditions and also the Wittgensteinian idea of communal agreement in usage as the furthest we can get by way of assigning such conditions.⁷ In which case there is no making sense of the claim – the ‘metaphysical’-realist claim – that a statement might be objectively true or false even though its truth-value lay beyond our utmost (whether present- or future-best) means of verification. Quite simply we have to regard such statements as concerning an ‘indeterminate’ area of reality and hence an area that offers no purchase for notions of objective truth or falsehood.

Other thinkers reject this full-strength version of the anti-realist argument since it conflicts so sharply with the rooted presumption that there *must* be a great many truths (mathematical, scientific, historical, and so forth) which lie beyond our furthest epistemic reach yet which can still be expressed in the form of hypotheses whose truth-value is determined independently of us and our limited knowledge.⁸ Thus there is something absurd about the notion that (say) mathematicians who sought a proof of Fermat’s Last Theorem before the proof was discovered, or physicists who hypothesised that ‘the charge on every electron is negative’ before that fact was established, must be thought of as making unintelligible statements or statements that were neither true nor false since they lacked adequate assertoric warrant. All the same such reservations may well go along with a high degree of epistemological scepticism or an outlook of principled agnosticism concerning the reality of objects and events that cannot be observed (or their existence verified) at first hand. Such is the verificationist refusal to entertain realist talk about entities (e.g. atoms or electrons) which play a crucial explanatory role in our current-best scientific theories but whose detection requires the use of advanced technological devices such as electron microscopes.⁹ This approach was adopted by the nineteenth-century physicist Ernst Mach who regarded such talk as a needless metaphysical extravagance and who advised that we should rather think of ‘atoms’ in strictly instrumentalist terms, i.e., as theoretical posits whose usefulness to science extended no further than their helping to account for certain observed (empirically verifiable) phenomena.¹⁰ A similar view was taken by the Vienna-Circle logical positivists, by the early Einstein who saw it as a means of countering commonsense-realist objections to Special Relativity, and more recently by Bas van Fraassen whose doctrine of ‘constructive empiricism’ likewise involves an unwillingness to credit the reality of anything that cannot be observed by the naked eye or without the aid of sophisticated means for enhancing or refining our perceptual capacities.¹¹ It is a version of the standard empiricist argument which holds that sensory acquaintance is the source of all knowledge – ‘nihil in intellectu quod non prius in sensu’ – and which eschews any recourse to theories or hypotheses concerning the existence of causal powers, dispositional properties, ‘hidden

variables' (in the quantum context), or other such merely notional appeals beyond the empirical evidence.¹²

Einstein famously renounced this Mach-inspired view when he confronted its full-scale programmatic statement in the form of the orthodox (Copenhagen) quantum theory. Thereafter, starting out from his series of debates with Niels Bohr, he insisted that the orthodox theory must be 'incomplete' in so far as it failed to deliver an intelligible realist and causal-explanatory account of the observed quantum phenomena.¹³ In more general terms the debate has most often been engaged between realists who typically defend some version of the argument for inference to the best explanation and anti-realists (or van Fraassen-type constructive empiricists) who insist that such arguments open the way to all kinds of epistemically unconstrained ontological licence.¹⁴ Moreover, this involves a sceptical assessment of the case for convergent realism since that case presupposes the existence (or reality) of at least a large proportion of the postulated objects – molecules, atoms, electrons, neutrinos, quarks, etc. – which the realist takes as sufficiently vouched by their explanatory (not just 'instrumental') role in our best, most successful theories to date. So there is plainly a sense in which 'constructive empiricism' shares a good deal with Dummett's more extreme metaphysical version of the anti-realist doctrine even though van Fraassen can present his thesis in the guise of a moderate (even 'commonsense') refusal to overstep the bounds of empirical adequacy and thereby yield unnecessary hostages to fortune.¹⁵ What it manifests, more specifically, is an attitude of deep-laid scepticism with regard to anything – whether causal powers or 'unobservable' entities – for which we are unable to claim direct epistemic or evidential warrant. This position has a clear intellectual ancestry in that phenomenalist idea of 'saving the [empirical] appearances' – without going so far as to assert any reality behind or beyond those appearances – which was first devised six centuries ago as a handy means of avoiding any clash between the new (heliocentric) astronomy and the dictates of theological orthodoxy. Such arguments were later revived by various thinkers – from Bishop Berkeley to Pierre Duhem – who likewise (albeit in different ways) sought to make room for religious faith by restricting the claims of scientific knowledge to the realm of empirical warrant.¹⁶

Of course these debates are nowadays conducted in a largely secularised context where such doctrinal pressures no longer exert much force. Nevertheless, that lineage is still apparent in the issues that divide realists from anti-realists, proponents of a causal-explanatory approach from empiricists of various type, and – at the most basic level – those who affirm the possibility of objective scientific knowledge from those who adopt (or profess to adopt) a standpoint of global scepticism. Above all, it comes out in the genuine bafflement that each party tends to evince when confronted with an advocate of the opposite view. Thus realists find it hard to conceive how anyone could possibly espouse a sceptical position given what they (the realists) take as the impressive record of achievement to date in the physical sciences and also the sheer self-evidence of those various propositions listed in my opening paragraph. To the sceptic, conversely,

that record is just a product of selective hindsight and that evidence no more than we are able to gather from our present-best methods of proof or information sources. Dummett puts the case most vividly when he contrasts the realist's willingness to accept that there are truths (or 'areas of reality') which may lie forever beyond our epistemic ken with the anti-realist's refusal to accept this predicament and hence their opting for a verificationist approach which makes such claims appear simply unintelligible. Thus:

[R]ealism about the past entails that there are numerous true propositions forever in principle unknowable. The effects of a past event may simply dissipate....To the realist, this is just part of the human condition; the anti-realist feels unknowability in principle to be intolerable and prefers to view our evidence for and memory of the past to be constitutive of it. For him, there cannot be a past fact no evidence for which exists to be discovered, because it is the existence of such evidence that would make it a fact, if it were one.¹⁷

Of course this argument must also apply to issues in epistemology and philosophy of science where it is likewise 'intolerable' – from the anti-realist's standpoint – to suppose that there might be unknowable truths about (e.g.) mathematics or the physical sciences whose objective standing is in no way affected by any limits on our best-available means of proof or verification. However, it will strike the realist as a flagrant example of the 'epistemic fallacy', that is, the kind of error that sceptics or anti-realists typically fall into when they confuse ontological issues with issues of knowledge or epistemic warrant.

This case has been developed most fully by Roy Bhaskar and other proponents of the Critical Realist approach to problems in epistemology and philosophy of the natural and social sciences.¹⁸ On their account, these problems arise through the conjunction of a narrowly procedural idea of scientific method with a failure to recognise the 'stratified' nature of reality, that is to say, the various kinds and degrees of human intervention in the process of acquiring scientific knowledge. Where realists sometimes go wrong – and lay themselves open to the standard range of anti-realist ripostes – is by treating issues of truth, rationality, empirical warrant, causal-explanatory scope, and so forth, as if they could somehow be resolved quite apart from the entire complex of physical and social factors that affect any given scientific experiment or the inferences drawn from it. In part this has to do with the lingering positivist tendency to adopt an idealised or sanitised view of experimental method which takes no account of such complicating factors or which assigns them to the background 'context of discovery' rather than the properly scientific 'context of justification'. Such approaches typically ignore what might be called the 'experimentalist's dilemma', i.e., the fact that it is only by conducting measurements or observations under highly artificial (laboratory) conditions that scientists can produce respectably uniform, precise, or law-like results. Thus, for instance, any set-up designed to quantify the effects of this or that physical force must involve a deci-

sion to disregard – or to factor out for observational purposes – a whole range of other interactive forces that operate both within and upon the given set-up. In which case there is a kind of implacable inverse-law by which the values of precision and generality are achieved only by abstracting away from the conditions of real-world applicability.

Of course scientists have little choice but to operate on this basis since their experiments would otherwise be subject to so many complicating factors and their results hedged around with so many qualifications as to undermine the very purpose of conducting such work. On the other hand – so Critical Realists argue – philosophers of science have less excuse for upholding a conception of scientific method that fails to make adequate allowance for them. And the best way to rectify this shortcoming is to take their point that reality is ‘stratified’, that is to say, composed of various ontological levels or domains, some of which are wholly independent of our sundry interests, methods, investigative procedures, techniques of observation, etc., while others are affected – more or less decisively – by just such forms of human intervention. Nor is this only a matter of coping with certain well-known problems in the realm of quantum mechanics where (at least on the orthodox interpretation) there is simply no appeal to an ‘objective’ quantum reality apart from the kinds of measurement carried out or the effect of observation on the system observed. To be sure, as I shall argue in subsequent chapters, such claims constitute a large challenge to any realist philosophy of science, and one to which realists have offered a range of philosophically cogent responses. However, what is distinctive about Critical Realism is its extension of arguments concerning the role of human interventionist agency to domains other than the microphysical, and indeed to *every* field of scientific enquiry where the methods and techniques of knowledge-acquisition are to some extent affected by physical, material, and social forces beyond any purely procedural (text-book) account. This approach brings various signal advantages, not least with regard to issues concerning the precise ontological status of certain items – such as transuranic elements or recombinant DNA proteins – that exist only in consequence of recent scientific and technological advances. Moreover, it allows for some highly productive thinking about the relationship between the physical and the social or human sciences, conceived as involving a complex, many-levelled process of exchange across boundaries that cannot simply be erased – as some postmodernists and Rortian ‘strong’-descriptivists would have it – but which are none the less open to constant renegotiation and critical dialogue.

Thus Critical Realism points a way beyond the old ‘two cultures’ debate, not to mention the latest outbreak of ‘science wars’ or the kinds of quarrel typically waged by philosophers who uphold the distinction between ‘context of discovery’ and ‘context of justification’ and sociologists of knowledge who reject that distinction *tout court*.¹⁹ This it does – to repeat – by avoiding any form of the epistemic fallacy that would confuse questions of objective truth with questions of humanly attainable knowledge, or ontological with epistemological issues. For it is just this confusion which has given sociologists a pretext for their claim that scientific knowledge is culture-relative or socially ‘constructed’ and their

outraged opponents a pretext for denouncing the errors and stupidities of people in the ‘soft’ (i.e., social-science) disciplines.²⁰ What Critical Realism thus holds out is the prospect of moving beyond such sterile debates through a more adequate conceptualisation of the various kinds and degrees of objectivity – or the various levels of interaction between knower and known – that distinguish not only the natural from the social sciences but also (say) physics from molecular biology, or genetics from embryology. All the same, it avoids any lingering commitment to the old hierarchical scheme of things according to which the disciplines could be ranked on a hard-to-soft scale of conceptual rigour or genuine scientific warrant. And this (I should add) without yielding any measure of hard-won realist ground to the sceptical or strong-descriptivist case for viewing such distinctions as merely a product of our present-day, culture-specific division of intellectual labour.²¹

II

It seems to me that Hume was right when he argued that scepticism could not be defeated on its chosen philosophical ground although it carried no weight against our settled convictions in every other context of enquiry.²² Where Hume went wrong was in thinking that these hyperinduced sceptical doubts were sure to arise once reflective individuals turned their minds to philosophy and were thus irresistibly driven to renounce all those tenets of the commonsense-realist outlook (like the existence of causal connections or laws of nature) which could not withstand such critical scrutiny. Indeed it is worth asking whether Hume, like Locke before him, would have pushed so far in this sceptical direction if the natural sciences of his day had offered the kinds of depth-explanatory grasp (e.g., the theories of chemical bonding or molecular and atomic structure) that have since lent considerable added support to the philosophic case for causal realism and inference to the best explanation.²³ Thus Locke might not have so forcefully denied the possibility of advancing from ‘nominal’ to ‘real’ essences or definitions had he been less struck by the limits of contemporary scientific knowledge with regard to just such matters. And again, scepticism concerning the reality of ‘unobservables’ like atoms was more rational for a nineteenth-century physicist-philosopher like Mach (at a time when their existence was strongly borne out on theoretical grounds but as yet far beyond the range of technologically assisted observation) than it is for a present-day thinker like van Fraassen (since they can now be observed and photographed). In short, the progress from instrumentalism to realism is one that tends to occur *pari passu* with advances in scientific knowledge and is therefore apt to undermine the sceptic’s case except where the sceptic constructs that case in just such a way as to deflect or invalidate any challenge from the realist quarter.²⁴

Nor can van Fraassen’s argument find much in the way of principled support from his notion that we don’t really *see* objects through electron microscopes, radio-telescopes, and other such complex devices since what we ‘see’ with their aid is perhaps just an artefact of the various technologies involved. For this is to

ignore three salient points: (1) that we can check those devices for accuracy, power of resolution, etc., through a range of well-proven test procedures; (2) that we possess sufficient understanding of their basic working principles to correct for certain errors; and (3) that it is the merest of anthropocentric delusions to suppose that such technologies are any less reliable – or any more subject to distorting effects – than the process of unaided human perception or observation with ‘the naked eye’. Moreover, it simply sets aside all the evidence from cognitive psychologists and neurophysiologists that what we ‘plainly observe’ is always already subject to a vast amount of interpretative processing by the visual cortex and related (though often topologically remote) regions of the brain.²⁵ Indeed, there is something distinctly odd about the way that van Fraassen combines such a high level of sophisticated philosophic argument with so naively restrictive and parochial an appeal to just those objects and events which happen to fall within the range of unaided human perceptual grasp.²⁶ Thus the realist need hardly be stuck for an answer when confronted with a case for ‘constructive empiricism’ which entails such a resolute refusal to acknowledge both the limits on our powers of ‘direct’ observation and the extent to which we can *explain and transcend* them through various kinds of applied scientific theory.

Of course there is deep resistance among some philosophers of science to the idea that epistemological issues could ever be resolved – or significantly advanced – through any discovery in the physical sciences. After all, what difference can it make whether scepticism is turned upon the kinds of claim that characterised previous stages in the history of science or on present-day claims such as those thrown up by the latest developments in quantum theory? Indeed if the realist takes that line then she is in for trouble with the well-known fact that quantum mechanics – at least on the orthodox interpretation – is fully as rife with conceptual problems and hence just as prone to scepticism as was ever the case with pre-Daltonian chemistry and physics.²⁷ At this point the realist may well respond that such problems are forced upon us *only if* we accept the orthodox veto on alternative theories – such as Bohm’s ‘hidden-variables’ account – which match all the known quantum predictive and empirical results while none the less providing a perfectly adequate causal-realist interpretation.²⁸ However, her general case against the sceptic is quite simply that things have moved on in the physical sciences and that arguments (like those of Locke and Hume) which once had a measure of rational-scientific (as well as philosophical) warrant must now look more like a rearguard defence of entrenched philosophical prejudice. Thus even the most doggedly orthodox of quantum theorists is working with a whole range of assumptions – concerning, for instance, the charge on certain particles, their interactive exchange of forces, or (most challenging to ‘classical’ realism) their wavelike distribution over a field of probabilistic measurement values – which implicitly deny the sceptical premise that there is no reality behind or beyond phenomenal appearances.²⁹ For whatever the problems in reconciling quantum mechanics with a realist worldview these are *epistemological* problems having to do with the scope and limits of our knowledge and not – as so often supposed – *ontological* problems that somehow

impugn the very notion of an objective and mind-independent quantum-physical reality. (Of course this claim needs a lot more in the way of detailed philosophical defence, for which I refer the reader to various passages of argument in chapters 3, 4, and 6.)

No doubt the sceptic can turn this argument around and insist that our present-best hypotheses – like so many in the past – will at length be consigned to the pre-history of some later, more ‘advanced’ physical theory, which will then in due course be abandoned with the advent of another (presumptively more adequate) account. Such – as we have seen – is the standard form of sceptical meta-induction or the ‘argument from error’ most often deployed by opponents of convergent realism. However, it ignores two main points which find strong support from detailed work in the history and philosophy of science. First, there is the fact that some past theories, though no longer accepted as scientifically adequate, can yet be seen to have approximated our current best knowledge in some degree and hence to retain a restricted validity or scope for application within certain specified limits. Such is famously the case with Newton’s theories of space and time and his account of gravitational attraction when interpreted within the more encompassing framework of Einstein’s theories of Special and General Relativity.³⁰ This argument extends, albeit more controversially, to the nineteenth-century notion of a ‘luminiferous ether’ which is often taken to have been refuted by the Michelson–Morley experiments but which – in the view of some historians of science – can be construed as making partial reference to a field-theoretical model like that described in Maxwell’s wave equations.³¹ Then again, there are instances – such as Black’s ‘caloric’ hypothesis to explain the properties of thermal conduction – which are now assumed to lack any reference to a strict ontological reckoning but which none the less played an important role in preparing the way for currently accepted theories, in this case the theory of specific heat.³² Thus it is wrong to conclude that *all* superseded hypotheses have completely dropped out of our scientific world-picture, like Aristotle’s theory of ‘natural place’ to explain why objects fall to earth when released or Priestley’s ‘phlogiston’-based theory of combustion. So the argument from error has no force against the argument for convergent realism just so long as one makes this crucial distinction between false (totally discredited) theories with empty referring terms and theories which can plausibly be represented as having pointed the way toward subsequent hypotheses that have so far stood up to the rigours of scientific testing.

On this account convergent realism and inference to the best explanation are themselves scientific hypotheses which should be treated pretty much on a par with the various causal-explanatory theories that they set out to justify.³³ Such is at any rate the realist case construed in epistemological terms, i.e., as it bears upon the prospects for our gaining reliable knowledge of the world through well-tried scientific methods and procedures. What scepticism amounts to in this context is a straightforward refusal – on philosophic grounds – to accept that such an argument can possibly provide conclusive evidence against the claim that all our experience is perfectly compatible with the non-existence of those various objects, properties, causal powers and so forth which we take (naively) as bearing

out the case for scientific realism. The best short answer to this is that the sceptic *can* indeed press his argument to the point of global philosophic doubt but that in so doing – in achieving such a pyrrhic victory – he will be left totally unable to account for our knowledge of the growth of scientific knowledge or the course of our everyday experience. Thus, as Michael Williams pithily puts it, '[t]he sceptic's fallacy is that he takes the discovery that, in the study, knowledge of the world is impossible for the discovery, in the study, that knowledge is impossible generally'.³⁴ Some philosophers – like Hume outside his study – recommend that we just get on with our lives since there is no answer to the sceptic (at least no cogent philosophical response) once our thoughts start running along this dead-end track. Others – like the later Wittgenstein – adopt a therapeutic approach and maintain that the sceptic's purported challenge simply doesn't make sense since there is no possibility of raising doubts with regard to such pointless or meaningless questions as that concerning the 'existence of an external world'.³⁵

However, this approach very often goes along with an appeal to those various practices, 'language-games', or communal 'forms of life' which are thought to offer all we need by way of reassurance but which the sceptic – or the cultural relativist – may readily turn to advantage. (For further discussion see Chapters 4 and 5.) That is to say, if 'agreement in judgement' – Wittgenstein's phrase – is the furthest we can get toward objectivity and truth, then such agreement will doubtless take as many forms as there exist diverse paradigms, worldviews, belief-systems, or conceptual schemes.³⁶ In which case Wittgensteinian 'realism' involves no substantive claim about the existence of objective (verification-transcendent) truths or of real-world objects, properties, or states of affairs which decide the truth-value of our statements concerning them. Rather – as exegetes like Saul Kripke have argued – it leaves no room for any but a 'sceptical solution' to the sceptic's challenge, i.e., a solution that falls back on communal warrant as our last, best resort in the face of such anxieties.³⁷ Thus, even as regards our most basic kinds of rule-following activity – addition, subtraction, continuing a series of natural numbers, and suchlike recursive operations – the sceptic can always ask: what guarantees that this is the *right* way to carry on rather than some alternative way that strikes us as plain wrong but which might just show that the subject in question is working with another (on its own terms perfectly consistent and intelligible) rule? For any standard that we might adduce for a proper application of the rule – like 'just keep adding two at every stage' – will then require a meta-rule for its own correct application, and this in turn a meta-meta-rule, and so on to the point of an infinite (vicious) regress.³⁸ Or again, the realist about arithmetical truths might think to block this regress by holding that the first-order rule is self-evident since it stands to reason for anyone who has grasped what is meant by 'addition', 'subtraction', 'continuing a number-series', etc. However – so the Kripkean sceptic rejoins – such arguments do nothing more than exchange a vicious regress for an equally vicious circularity of reasoning which fails to meet the sceptical challenge on precisely the main point at issue. For, according to Kripke, there is a close connection between Wittgenstein's thoughts about rule-following and those other passages

in the *Philosophical Investigations* where he rejects the idea of a ‘private language’, that is to say, an inner realm of thoughts, meanings, or intentions to which the solitary thinker has privileged epistemic access and which thus puts a stop to any doubts concerning what they have in mind when following this or that rule.³⁹ Such is the still prevalent Cartesian illusion (as Wittgenstein sees it) that we can somehow check the correctness of our concepts and meanings by comparing them with – what else? – our concepts and meanings, as if we were to buy a second copy of the daily newspaper just to make sure that the first copy contained a fully accurate report of events.

So it is a fallacy to think that the rule-following paradox could ever be resolved by an appeal to some putative ‘fact’ about what competent reasoners have in mind when they take themselves to be correctly performing some basic arithmetical or logical operation. For, quite simply, there is no such fact to be known and hence no appeal to a standard of correctness that would enable us to say (with total conviction) that a maths student who was asked to add $68 + 57$ and replied ‘125’ was correctly applying the rule of addition whereas another student who was posed the same question and replied ‘5’ or ‘392’ either hadn’t grasped the operative rule or – for some reason – must have failed to apply it in this particular case. Thus, the second student could always, with sufficient ingenuity, come up with some alternative (perhaps highly complex and *to us* weirdly counter-intuitive) rule which made his answer right and the first student’s answer wrong. At any rate there is nothing that can block the path to this sceptical conclusion if one takes it – on the view that Kripke is here attacking – that correctness in such matters comes down to a question of applying the rule in accordance with principles that are somehow self-evident to reason or which somehow correspond to our inward grasp of what constitutes a valid application. And indeed the sceptical argument is sure to go through if the issue is set up in these terms, i.e., as a straight choice between supposing (deludedly) that our knowledge of the rules can be checked against our knowledge of the rules and concluding rather – as Kripke would have it – that the criteria for correctness cannot be other than those supplied by some communal practice or relevant (e.g., arithmetical) ‘form of life’. For the circularity (or the vicious regress) will always reappear at some point so long as the realist can be forced or persuaded to accept this ‘Kripkensteinian’ dilemma.

However, there is a strong case for maintaining that in fact it is a false dilemma and one that could scarcely have posed such a challenge – or provoked such a range of conflicting responses – were it not for the current high standing of Wittgenstein’s later philosophy. This has the effect of convincing many thinkers (among them realists of various technical persuasion) that their arguments will need to go by way of an encounter with Wittgenstein on ‘private languages’ and ‘following a rule’ if they are to have the least chance of holding up against sceptical attack. Thus – to take perhaps the most prominent example – Hilary Putnam’s work over the past four decades in epistemology, philosophical semantics, and philosophy of science and mathematics can be seen to have moved very largely under Wittgenstein’s influence from a full-fledged realist position to

various forms of 'internal', framework-relative, or (in his most recent writings) 'naturalised' quasi-realism.⁴⁰ Elsewhere, among philosophers like Thomas Nagel and John McDowell, one finds a similar compulsion to run their arguments through the Wittgensteinian hoop even though a great deal of what they have to say – in Nagel's case especially – comes out sharply opposed to the idea that mathematical truths or the findings of the physical sciences could ever be accounted for in terms of our various communal practices or shared 'forms of life'.⁴¹

Where this compulsion most often gets a hold, I think, is through Kripke's sharpening of the Wittgensteinian pseudo-dilemma to the point where it seems that there is no alternative except to endorse *either* his sceptical (communitarian) 'solution' to the sceptical paradox *or* some version of the hopeless (regressive or circular) appeal to standards of veridical warrant 'in the mind' of this or that reasoner. However, it is at just this crucial point that the well-advised realist will refuse to accept the terms on offer and insist that they involve – like so many other versions of the sceptical or anti-realist case – a plain confusion between ontological and epistemological issues. That is to say, what the Kripkensteinian 'paradox' preemptively excludes is the claim that statements such as ' $68 + 57 = 125$ ' or 'the charge on every electron is negative' have their truth-value fixed by the way things stand in mathematical or subatomic reality, rather than through their accordance *either* with our 'private' (inward) state of conviction regarding such matters *or* with whatever practices and beliefs hold sway within some given community. Thus Kripke's knockdown challenge to realism can be seen to trade on a very odd notion of what 'realism' entails, that is, the idea – despite his avowedly anti-Cartesian line of argument – that to be any kind of realist about (say) the truths of elementary arithmetic is to espouse some version of the argument from privileged epistemic access and hence to fall plump into Wittgenstein's 'private language' trap. But this challenge will scarcely impress the realist who keeps a firm grip on the basic distinction between *certainty* and *truth*, or (more generally) those states of mind that characterise our various degrees of epistemic warrant and those real-world states of affairs – or objective truths – which decide whether or not we are justified in holding such-and-such to be truly, necessarily, probably, or possibly the case.

Scott Soames makes this point to telling effect when he remarks that the truth-value of the statement 'there is 0.5 probability of the tossed coin's turning up heads' also holds good for the statement 'the probability that this is true = 0.5' but *not* for the statement 'it is 0.5 certain that the coin will turn up heads'.⁴² What gives the Kripkensteinian 'paradox' its seeming force is the idea that talk of 'objectivity' or 'truth' can always be construed as involving some appeal to certainty, and that any such appeal (since it rests on the notion of privileged or 'private' access) must therefore lie open to the standard sceptical rejoinder. But this is to ignore not only, as I have said, the distinction between ontological and epistemological issues but also that between certainty (or the lack of it) as a matter of objective statistical warrant and certainty (or *our* lack of it) as a matter of epistemic or predictive warrant. After all, as Soames very pointedly asks, '[w]hy should the claim that certainty is unattainable lead one to think predications of

truth or falsehood have no place in actual science?’⁴³ Thus we risk all sorts of confusion with regard to the realism/anti-realism issue if we fail to distinguish the realm of objective truths (including objective probabilities) from the realm of epistemological enquiry where such truths are always what we are aiming to know but where our present-best state of knowledge may always fall short in some respect or degree.

III

I have suggested here (and will argue more fully in Chapter 4) that the grip of these confusions on so much recent philosophical thought can be attributed in part to Wittgenstein’s influence and its effect in promoting a verificationist approach to issues of knowledge and truth. More specifically, it derives from the Fregean principle that ‘sense determines reference’ joined to the Wittgensteinian precept that the sense (or operative meaning) of our statements is a matter of their verifiability-conditions as given by their usage in various recognisable contexts of utterance. Thus – as Dummett saw it in 1978 – ‘[t]he whole point of my approach...has been to show that the theory of meaning underlies metaphysics. If I have made any worthwhile contribution to philosophy, I think it must lie in having raised the issue in these terms.’⁴⁴ And this view had not changed by the time of his valedictory lecture at Oxford in 1993 when he remarked:

[t]he opinion is sometimes expressed that I succeeded in opening up a genuine philosophical problem, or range of problems, but that the resulting topic has little to do with traditional disputes concerning realism. That was certainly not my intention: I meant to apply a new technique to such wholly traditional questions as realism about the external world and about the mental, questions which I continue to believe I characterised correctly.⁴⁵

However, the main issue is not so much this exegetical point as to whether we have got Dummett’s precise intentions right or wrong but rather the question whether Dummett is himself right or wrong in treating ‘metaphysical’ issues about truth and objectivity as best approached through a theory of meaning that precludes any legitimate appeal to the existence of objective (verification-transcendent) truth-values. Not that Dummett goes so far as a sceptic like Kripke in denying that even sentences with well-defined assertibility conditions are apt for judgements of truth or falsehood. Where Kripke takes this to follow directly from the Wittgensteinian ‘paradox’ about rule-following, Dummett adopts the less extreme view that sentences possessing such conditions can harmlessly be treated as fit candidates for assessment in these terms. However, he still very firmly maintains – with presumptive warrant from Frege and Wittgenstein – that if meaning is given by truth-conditions and if truth-conditions depend on our capacity to acquire and to recognise the criteria for correct usage then correctness must itself come down to a matter of shared or communal warrant.

So there is less difference than might at first appear between Kripke's far-out sceptical take on these lessons from Wittgenstein and Dummett's attempt to follow them through without, in the process, giving rise to so starkly paradoxical an upshot. In Dummett this results from his belief that any statement of the case for realism must go by way of a theory of meaning according to which we *can* indeed acquire and recognise the truth-conditions for undecidable sentences or those that belong to the so-called 'disputed class'. Which is also to say that the dispute between realists and anti-realists is first and foremost a dispute within the realm of philosophical semantics and only then (once those issues have been resolved) a dispute concerning other, more 'traditional', i.e., metaphysical and ontological, questions. However, the realist is likely to object that Dummett's way of setting the agenda is one that effectively decides the outcome in advance since it preempts or excludes the alternative view that such questions have *nothing whatever to do* with any theory of meaning, whether 'realist' or 'anti-realist'. Michael Devitt puts this case in typically forthright and (to my mind) convincing terms. 'Realism', he writes,

is an overarching empirical (scientific) theory or principle. It is initially plausible. It is supported by arguments that make no appeal to theories of language or understanding...What firmer place could there be to stand than Realism, as we theorize in such undeveloped areas as those of language and understanding? In contrast, the poor state of theories in those areas, whether verificationist or not, makes them a bad place from which to start theorizing, particularly in determining overarching principles about the nature of reality. To think otherwise is to put the cart before the horse.⁴⁶

Thus the trouble with Dummett-style anti-realism is that it starts out from a stipulative notion of what the adversary case *must* amount to – namely, a realist *theory of meaning* which allows for the existence of verification-transcendent truths – if that case is to carry any philosophic weight with (who else?) the anti-realist. Whence it proceeds to demonstrate, sure enough, that such a theory cannot be had since *ex hypothesi* any statement concerning truths which transcended our utmost powers of verification would be a statement whose truth-conditions were beyond our recognitional (or linguistically manifestable) grasp. So the realist is left in a no-win situation, one that constrains her to formulate her case in terms of a theory which prevents that case from even getting off the ground.

Devitt provides the best short response when he asks: 'What has truth to do with Realism?' and answers: 'On the face of it, nothing at all.' And again, '[r]ealism says *nothing semantic at all* beyond, in its use of "objective", making the negative point that our semantic capacities do *not* constitute the world'.⁴⁷ To Devitt's way of thinking anti-realism can be made to look plausible only through this regular confusion between issues of truth (nowadays conceived in linguistic or logico-semantic terms) and issues concerning the nature, structure, and properties of physical reality. Thus it is (one might say) a strictly *preposterous* doctrine in the etymological sense of that term, a doctrine that 'puts the cart before the horse' by supposing that first-order questions of this kind could ever be resolved – or ruled out of court – by some

second-order theory with shaky credentials in post-Wittgensteinian philosophy of language. It is worth quoting Devitt at greater length since he brings out exactly what is wrong (from a realist viewpoint) with this whole recent tendency to shift the emphasis from ontological to epistemological issues and thence to debates in the theory of meaning. Such arguments typically start out, he remarks,

with a properly metaphysical statement of the issue. This is immediately replaced by a formulation in terms of truth, which is then taken...as part of a theory of meaning. Whatever the merits of the various theories of meaning then proposed, the theories are (almost) irrelevant...to the metaphysical issue which they are alleged to settle. For the metaphysical issue is not one about meaning.⁴⁸

Of course there are alternative (realist) conceptions of truth which don't thus reduce it to a role within the theory of meaning and thence – *via* the doctrine of meaning-as-use – to a matter of 'recognition' or 'manifestation' by subjects suitably placed to judge of its applicability in this or that context. Indeed, one might ask whether Devitt's radical disjunction between 'realism' and 'truth' doesn't on occasion risk depriving his realist metaphysics of any adequate grounding in epistemology or in those various truth-conducive procedures – such as causal reasoning or inference to the best explanation – that have characterised the growth of scientific knowledge.⁴⁹ Such is at any rate the implication when Devitt asserts that truth has 'nothing at all' to do with realism and that this whole debate has been skewed from the outset by those (like Dummett) who persist in confusing these two quite distinct issues.

Still it is clear that Devitt finds room enough for other, more congenial ways of restoring the link between them. One is the modal-realist approach developed by Kripke and the early Putnam as part of their case for a causal theory of reference that would avoid certain problems (such as radical meaning-variance between different theories or conceptual schemes) which beset the Frege–Russell descriptivist account.⁵⁰ In my view – as in Devitt's – this causal theory has by far the best claim to make sense of our knowledge of the growth of scientific knowledge, whatever Putnam's subsequent doubts and whatever the impact on recent debate of Kripke's so-called 'sceptical solution' to the rule-following paradox. Thus it aims to re-establish fixity of reference for natural-kind terms like 'water' and 'gold' by supposing that those items were first picked out by a designative act of naming ('this is water', 'this is gold') which thereafter continued to denote the same sorts of stuff despite various – sometimes radical – shifts in our range of scientifically informed identifying criteria.⁵¹ Thus *water* went from something like 'transparent, odourless, (normally) liquid substance that freezes or boils at certain temperatures, falls as rain, fills up lakes, possesses certain cleansing properties', etc., to 'substance with the distinctive molecular structure H₂O'. And *gold* was subject to a similar process of increasing definitional refinement from 'yellow, ductile, metallic substance that dissolves in dilute nitric acid' to 'metallic element with atomic number 79'. So likewise with *acid* where the change went roughly

from ‘corrosive substance which reacts in certain ways when mixed with other kinds of stuff’, to ‘substance which turns litmus-paper red’, and thence to our present understanding of it as ‘having the property of proton-donor’. Such examples can readily be multiplied across the whole gamut of the natural sciences, so that (for instance) what properly counts as a *lemon* was once ‘bitter-tasting fruit of a certain shape with a yellow skin and white rind’ whereas now we can pick out lemons more reliably – and withhold that name from fruits which look very much like lemons but in fact belong to a different species – by reference to their chromosomal structure. (Also, of course, we can identify as a genuine *lemon* any unripe [green] sample that happens to be saturated with sugar and hence sweet to the taste.) And in the case of *tiger* we now know enough about the genetic constitution of tigerhood to settle the issue with regard to what counts as a member of that species as distinct from some other kind of creature that looks and behaves very much like a tiger (striped, fleet of foot, carnivorous, of large-scale feline appearance, etc.).

It is the chief virtue of this theory – according to Kripke and early Putnam – that it manages to explain how reference is fixed, how it remains sufficiently stable through episodes of subsequent theory-change, and also how early usages of a term are ‘truth-tracking’ or ‘sensitive to future discovery’ in so far as water *always was* H₂O and gold *always was* the metallic element with atomic number 79 even when these truths were unknown and unknowable by any existing scientific means.⁵² Moreover, they are taken as belonging to the class of *a posteriori* necessary truths, that is, as having resulted from certain empirical discoveries about the physical world and its microstructural constitution but nevertheless as holding *necessarily* for any world that resembles our own in the relevant respects, such as that of containing certain elements, molecules, chemical-bonding properties, genetic-chromosomal structures, and so forth. Hence Putnam’s famous series of ‘Twin-Earth’ thought-experiments where he asks us to imagine another planet very much like ours but where what they call ‘water’ has the molecular composition XYZ instead of H₂O, where what they call ‘aluminium’ has the atomic structure of Earthian molybdenum and vice versa, or where the creatures which they (the denizens of Twin-Earth) call ‘tigers’ turn out to be silicon-based rather than carbon-based life-forms.⁵³ In each case, so Putnam persuasively argues, the reference is fixed (and proper usage effectively constrained) by whatever it is – at this microstructural level – that Earthians or Twin-Earthians pick out when they travel to the counterpart planet and think to identify ‘water’, ‘gold’, ‘aluminium’, ‘tigers’, and the rest. Thus the Earthian space-travellers are wrong – misled by phenomenal appearances – if they suppose Twin-Earth to contain large quantities of *water* (which for them of course refers to H₂O), when in fact what they behold is a planet plentifully stocked with XYZ. And the visitors from Twin-Earth are likewise wrong if they report ‘Lots of tigers on Earth and all the saucepans are made of aluminium, just like ours’ when in fact the Earthian tigers only superficially resemble their Twin-Earth (silicon-based) ‘tigers’ and the saucepans are indeed made from *aluminium* (the metal with just that atomic structure) but are

wrongly so described by the visitors because their ‘aluminium’ is in fact Earthian *molybdenum*.

Outside the realm of science-fiction fantasy this argument would apply in modified form to real-world historical instances such as the mistaking of iron pyrites (or ‘fool’s gold’) for samples of the genuine kind, or various pre-Daltonian confusions with regard to elements and compounds, or the mistaking of whales for a species of fish. These cases differ from the Twin-Earth examples in so far as they involve an inadequate (primitive or superficial) grasp of just what it is that constitutes the kind in question, rather than a perfectly adequate grasp of constitutive features that happen to obtain in another (differently constituted) world. However, this is an additional merit of the early-Putnam approach, i.e., that it affords a wide range of distinctions between theories (like Priestley’s phlogiston-based account of combustion) which turned out to be false since they contained empty or non-referring terms, theories (like those of the atomists from Democritus to Dalton) which were underdeveloped in various degrees though ultimately on the right track, and theories – like the present-day standard model of subatomic structure – which, if true, are *necessarily* true in virtue of the way things stand in physical reality. Thus, it also allows (as the realist surely must) for the non-finality of scientific knowledge at any given stage of enquiry and the strong likelihood that many of our currently accepted theories will end up at best among those that were on the right track and at worst among those that were on the wrong track entirely. However, as I have said, this is no argument against scientific realism since it is just the realist’s cardinal point – and the chief point of dispute by anti-realists – that truth can always in principle transcend the limits of our present-best knowledge or means of verification. In short, what the Kripke–Putnam theory of reference provides is a highly developed philosophical account which *does* find room for a close relation between issues of realism and issues of truth but which *doesn’t* equate the truth-aptness of our various statements with our capacity to verify those statements or to manifest a grasp of their truth-conditions by exercising just that capacity. Thus there is no reason – metaphysical prejudice aside – to suppose (with Dummett) that realism must stand or fall on its ability to meet the conditions laid down by a verificationist theory of meaning.

Indeed it might be said that my entire discussion so far, like Devitt’s response to Dummett, provides just another cautionary instance of the lengths to which realists are reluctantly forced in their attempt to meet the challenge of an anti-realist position which simply misses the main point about realism. Hence perhaps the somewhat irritable tone that enters Devitt’s argument when he remarks that semantic theories of truth are ‘almost irrelevant...to the metaphysical issue which they are alleged to settle’, or that realism ‘says nothing semantic at all’ aside from ‘making the negative point that our semantic theories do *not* constitute the world’.⁵⁴ The same applies to those various responses to Kripke on the rule-following ‘paradox’ that often go some highly elaborate ways around in defending the claim that arithmetical truths are objective and in no sense a matter of ‘agreement in judgement’ or conformity with communal practice. For

there is – to put it plainly – something absurd about a theory which purports on semantic grounds to show that we have no better warrant than this for such statements as ‘ $2 + 2 = 4$ ’ or ‘ $65 + 58 \neq 7$ ’. And there is likewise (I submit) something absurd about the claim, whatever its presumptive authority from Wittgenstein, that well-formed but undecidable statements such as ‘Goldbach’s Conjecture is true’ or ‘there exists a solar system like ours in some radio-telescopically inaccessible region of the expanding universe’ cannot possess an objective truth-value just because we are not yet (or may never be) in a position to confirm or disconfirm them.⁵⁵

IV

The latter sort of case – from astronomy – is one that Dummett invokes metaphorically by way of characterising the realist (or ‘Platonist’) conception of mathematical truth. On this account, he writes, ‘mathematical structures, like galaxies, exist, independently of us, in a realm of reality which we do not inhabit but which those of us who have the skill are capable of observing and reporting on’.⁵⁶ Dummett of course rejects this idea in favour of an intuitionist conception which holds mathematical truth within the limits of whatever we are able to prove or verify by the best formal means at our disposal.⁵⁷ More than that, he argues: it is an error to suppose that any statement concerning the ‘structures’ in question could possess an objective truth-value aside from our ability to manifest a grasp of their operative truth-conditions. Rather, we should take the more cautious (verificationist) line of allowing that statements about galaxies and mathematical structures are candidates for warranted assertibility just so long as they meet those criteria while otherwise – as in the case of unobservable galaxies or unproven theorems – they belong to the ‘disputed class’ and are hence ruled out for such treatment. Thus truth-values have no application to whatever lies beyond our epistemic ken and apply only to whatever we can demonstrably claim to know or observe. In both cases – galaxies and theorems – our statements are altogether void of truth-content if they lie beyond the range of verification and should be taken to possess such content only in so far as they meet the conditions for warranted assertibility. From which it follows that the galaxies, like the theorems, are structures whose very existence must finally depend on our capacity to acquire or to manifest knowledge of them. Such is the logic of Dummett’s anti-realist argument – his idea that any ‘gaps in our knowledge’ must also be construed as ‘gaps in reality’ – despite the occasional signs in his later work that he is willing to accept a more qualified version of the thesis.⁵⁸

As I have said, there is something decidedly *outré* about the notion that galaxies – even (or especially) those in some remote and inaccessible region of the expanding universe – should be thought of as subject to the limiting conditions of human epistemic grasp. In the case of mathematics, anti-realism looks more plausible since it seems to capture the sheer impossibility of conceiving that there are certain abstract and mind-independent objects (numbers, sets, classes, functions, and so forth) to which none the less we can gain access through

some mysterious, quasi-perceptual mode of apprehension.⁵⁹ According to Dummett this is a dilemma that will always confront the mathematical realist when she seeks to explain how and where those putative realia can be thought to exist or to have their dwelling. In other words it is a version of the standard argument (from Aristotle down) against Platonism of whatever kind: namely, that there is just no point in appealing to a realm of supra-sensory ‘forms’ or ‘essences’ which by very definition transcend the limits of human epistemic grasp. Hence the pyrrhic conclusion of some philosophers, middle-period Putnam among them, that quite simply ‘nothing works’ in philosophy of mathematics since we can *either* have knowledge (in which case objectivity goes by the board) *or else* hang onto the notion of objectivity (in which case mathematical truths are forever unknowable).⁶⁰ However, this looks like yet another forced dilemma of the kind that anti-realists very often exploit in order to press their arguments home. For it is only on a certain (*echt*-Platonist) conception of what mathematical realism entails that its proponent can be made out to endorse the plainly self-contradictory idea of our somehow having quasi-perceptual acquaintance with abstract or supra-sensory truths.

Jerrold Katz nicely captures the absurdity of this notion in a passage that is worth quoting at length since it also puts the case for an alternative realist approach in philosophy of mathematics that draws on some of the modal distinctions that I have discussed in the course of this chapter. Thus:

[t]he entire idea that our knowledge of abstract objects might be based on perceptual contact is misguided, since, even if we had contact with abstract objects, the information we could obtain from such contact wouldn’t help us in trying to justify our beliefs about them. The epistemological function of perceptual contact is to provide information about which possibilities are actualities. Perceptual contact thus has a point in the case of empirical propositions. Because natural objects can be otherwise than they actually are (*non obstante* their essential properties), contact is necessary in order to discover how they actually are....Not so with abstract objects. They could not be otherwise than they are....Hence there is no question of which mathematical possibilities are actual possibilities. In virtue of being a perfect number, six must be a perfect number; in virtue of being the only even prime, two must be the only even prime. Since the epistemic role of contact is to provide us with the information needed to select among the different ways something might be, and since perceptual contact cannot provide information about how something must be, contact has no point in relation to abstract objects. It cannot ground beliefs about them.⁶¹

What is so odd about Putnam’s counsel of despair – his idea that ultimately ‘nothing works’ in philosophy of mathematics – is the fact that his own earlier work (along with Kripke’s) offered such a range of promising resources for avoiding such a hard-put sceptical upshot. That is to say, it distinguished clearly between contingent matters of fact about the world, *a posteriori* necessary truths

(e.g., concerning subatomic or molecular structures, chemical properties, or laws of nature) which hold across all possible worlds that are physically congruent with ours, and *a priori* truths – such as those of mathematics – whose necessity follows from just *what it is* to be a prime or a perfect number. Above all it shows the fallacy contained in anti-realist arguments to the general effect that anyone who takes a realist (or objectivist) view of mathematical truth must also be a Platonist in the full-fledged sense of supposing that our epistemic access to such ‘things’ as numbers, sets, or classes involves some wholly mysterious kind of epistemic ‘contact’ with them. What such arguments ignore – for their own sceptical purpose – is precisely the point, as Katz remarks, that perceptual contact has the role of providing us with information about things (objects and events) that might have been otherwise, and therefore has no role whatsoever in our dealing with abstract entities. So the realist’s supposed no-win predicament – that we can *either* have mathematical truth *or* mathematical knowledge but surely not both – is one that is surreptitiously foisted upon her by a false conception of what realism amounts to in the mathematical domain.

It is this pseudo-paradox that drives a great many current anti-realist approaches, among them Dummett’s philosophy of mathematics, Kripke’s so-called ‘sceptical solution’ to Wittgenstein’s problem about rule-following, and – not least – Putnam’s later idea that truth must be relative or ‘internal’ to some given investigative framework or conceptual scheme, whether in the physical sciences or in the formal disciplines of logic, set-theory, and even elementary arithmetic.⁶² Not that the issue is likely to be settled or the sceptic at last won over by any argument that involves (as at some point it must) an appeal to just those realist principles which he is out to challenge. Thus the hard-line sceptic concerning these matters can always stick to his guns and deny that we have any reason for accepting inference to the best causal explanation or the necessity that ‘ $68 + 57 = 125$ ’ is a true statement whatever the scope for alternative ‘correct’ solutions according to some different rule. Then again he might profess to doubt the validity of even such axiomatic truths as Newton’s First Law of Motion, i.e., the inertial law that a body will continue in a given state of motion unless and until acted upon by some external force, along with all its jointly mathematical and physical consequences. However, in that case the sceptic would be pressing his argument to a point where it rendered all of physics unintelligible (including our very conception of a physical body) since it failed absolutely to explain why mathematics should have played so crucial a role in the advancement of the physical sciences from Galileo down.

One is tempted to conclude, like Kant about Hume, that if scepticism of this sort is the inevitable upshot of philosophic reflection on the problem of knowledge then that fact is more of a scandal to philosophy than a cause for anxiety among scientists or layfolk.⁶³ Kant’s answer to Hume – his doctrine of transcendental idealism conjoined with empirical realism – has scarcely managed to avert that scandal despite his confident claim to have provided the wished-for deliverance from all our sceptical doubts. Indeed it may be argued that Kant’s ‘solution’ has since given rise to most of the problems currently exploited by sceptics,

constructivists, and anti-realists of various persuasion. However, as I have argued here, philosophy of science now has a range of more adequate resources with which to counter that challenge, among them the modal-realist conception of *a posteriori* necessary truths, which found no place in Kant's metaphysics. Indeed it was precisely for lack of such resources that Kant was on the one hand led to construe causality in mind-dependent terms (thus laying himself open to the charge of downright idealism) and on the other hand led to claim absolute *a priori* warrant for the axioms of Euclidean geometry (which were later shown to constitute just one among the various possible or geometrically consistent systems).⁶⁴ Nor is it any coincidence that a thinker like Putnam has retreated from his earlier realist position very largely in consequence of the sceptical reflections prompted by this and other problematic episodes in the history of post-Kantian scientific thought.⁶⁵ Thus there is good reason to conclude that modal realism is the only approach in philosophy of science and mathematics that affords an adequate line of defence against scepticism and anti-realism in their sundry present-day forms. To this extent, moreover, it is the sole approach that succeeds in avoiding the philosophic 'scandal' of a doctrine which denies both the very possibility of objective mathematical and scientific truths and the prospect of our coming to know at least some of them through various (albeit fallible) methods of enquiry.

Such is at any rate the best case for scientific realism as viewed by some of its leading present-day advocates. What I propose to do now – in chapters 2 to 5 – is examine how it fares under sceptical pressure from an exceptionally shrewd and well-informed philosopher-historian of science, Norwood Russell Hanson, whose thinking has exerted a powerful (though often unacknowledged) influence on debates over the past half-century. In chapters 6 and 7 I shall come back to some of the issues raised above with particular reference to quantum mechanics – since it figures so centrally in Hanson's thinking – and to present-day arguments in epistemology and philosophy of science.