

Part II

Current Research and Issues

A. General Issues in Philosophy of Science

5 Scientific Realism with a Humean Face*

Stathis Psillos

1. Introduction

This chapter offers an intellectual history of the scientific realism debate during the twentieth century. The telling of the tale will explain the philosophical significance and the prospects of the scientific realism debate, through the major turns it went through. The emphasis will be on the relations between empiricism and scientific realism and on the swing from metaphysics-hostile to metaphysics-friendly versions of realism.

2. From Verification to Confirmational Holism

The early stages of the debate were shaped by the verificationist criterion of meaning (VCM). This criterion brings together semantics (issues about meaning) with epistemology (issues about knowledge): meaningfulness is tied to verifiability and meaning to verification. The key idea is that the only non-analytic statements which are meaningful are those whose truth can be established empirically; concomitantly, *the* meaning of a synthetic statement is its empirical content, *viz.*, whatever part of its content can be established empirically and *no more*. Verificationism was the principal way to capture the deep anti-metaphysical commitments of logical positivism. It was a way to show that there was no real excess content to metaphysical statements that were supposed go beyond what is verifiable in experience. A semantic criterion with a distinctively epistemological dressing was also supposed to be a criterion for demarcating science from metaphysics.

In Moritz Schlick's (1932) hands, VCM was meant to deliver science from metaphysics, *without*, however, revising the rich conception of the world, as this is described by the sciences. Schlick was quite adamant that VCM, in its anti-metaphysics capacity, was leaving the world as described by science intact – a world populated by atoms and fields and whatever else our best science tells us there is – this is what he called 'empirical realism'. VCM also dictated a certain

solution to the problem of empirically equivalent descriptions of the world. It may well be that the issue between realism and idealism is a pseudo-problem, since VCM licences no empirical difference between the two (see Carnap circa 1928), but – by the very same token – the issue between competing but empirically equivalent *scientific* theories of the world becomes a pseudo-problem, too. Take, for instance, the rivalry between the general theory of relativity, which has it that the structure of space is non-Euclidean, and a version of Newton's theory, which keeps Euclidean geometry and posits universal forces, which acts indiscriminately on all moving bodies (e.g. moving rods) and makes them to contract accordingly. These two theories, as Reichenbach showed, are empirically indistinguishable. Hence, according to VCM, the very choice between them would end up being a pseudo-problem. The natural reaction to this problem was to claim that, in line with VCM, empirically indistinguishable theories are cognitively equivalent – different but ultimately equivalent (inter-translatable) formulations of the same theory. If rival theories, *properly understood*, end up being the same theory, it is no longer the case that VCM leaves the world as described by science intact. (This, however, was not Reichenbach's reaction, who adopted a *probability theory of meaning* precisely in order to break such deadlocks issued by VCM.)

In any case, verificationist semantics cannot really capture the fine structure of the relations of confirmation between evidence and theory. The challenge here is double. The first comes from the fact that the relations between evidence and theory are probabilistic and not deductive (as was pointed out by Hans Reichenbach 1938). The other comes from the fact that even if we were to stick to a strictly deductive account of confirmation, theoretical assertions always have excess content over their observational consequences (as was realized by Rudolf Carnap in his 1937a). The problem that Carnap uncovered in *Testability and Meaning* was precisely that insofar as verificationism is uncoupled from reductive versions of empiricism, semantics had to be liberalized in such a way that verifiability gives way to a weaker notion of confirmability.

The move from verification to confirmation was groundbreaking. The ensuing liberalization of empiricism didn't challenge the subordination of metaphysics to semantics. But, in the longer run, it did pave the way for taking metaphysics seriously. Carnap favoured the 'requirement of confirmability': every synthetic statement must be confirmable (CT). This, of course, is an extremely liberal criterion which renders meaningful all kinds of theoretical assertions, as well as nomological statements. But Carnap did not think that the anti-metaphysical stance of empiricism (so far served by VCM) was thereby abandoned. The new liberalized criterion of meaningfulness (CT), Carnap (1937a, p. 35) says, 'suffices to exclude all sentences of a non-empirical nature, e.g., those of transcendental metaphysics inasmuch as they are not confirmable, not even incompletely.'

This declaration was more like wishful thinking. The problem lied in a fact that Carnap himself was fully aware of, viz., that confirmation is holistic. With due acknowledgement to Duhem and Poincaré, he noted already in 1937:

It is, in general, impossible to test even a single hypothetical sentence. In the case of a single sentence of this kind, there are in general no L-consequences of the form of protocol sentences; hence for the deduction of sentences having the form of protocol-sentences the remaining hypotheses must also be used. Thus the test applies, at bottom, not to a single hypothesis but to the whole system of physics as a system of hypotheses (Duhem, Poincaré). (1937b, p. 318)

Hence, there is a tension in the offing. Confirmation-based semantics does leave the world as described by science intact, but it now seems that metaphysical assertions might not be sharply separable from scientific ones. Almost concurrently with Quine's famous attack on analyticity and semantic atomism (the two dogmas of empiricism), Hempel (1951) showed in some detail that empiricists had failed to formulate general and precise criteria which can separate some (metaphysical) statements as isolated without rendering other meaningful statements isolated, too. Despite Carnap's heroic efforts to formulate a criterion by means of which metaphysical statements would end up being isolated, while proper scientific ones, no matter how remote from experience, would not, he failed (see Psillos 2008, for the details).

By the middle of 1950s, there was full recognition among empiricists of the meaningfulness of theories and of the excess content that theoretical assertions have over their observational consequences. Yet, the very movement of thought that led to this conclusion had put in jeopardy the standard empiricist attempt to sharply separate science from metaphysics.

3. The Battle of Empiricism – Phase I

Here starts one of the most interesting battles within empiricism. A significant new addition was Herbert Feigl (1950), whose main point was that once it was accepted that theoretical terms have 'excess content', and once VCM is abandoned, it is but a short and harmless step to accept that theoretical terms have factual reference: they designate theoretical/unobservable entities. The so-called surplus meaning of theoretical terms – whatever is not characterized in terms of their observational consequences and their links with other theoretical terms – consists in their factual reference, where 'in the language of empirical science all those terms (and only those terms) have factual reference which are linked to each other and to the evidential basis by nomological relationships' (1950, p. 50).

This suggestion may not seem enough to guide ontological commitment. But this is not true. For one, it makes it clear that unobservable entities are no less real than observable entities, given that, as Feigl put it, 'they are on a par within the nomological framework' of modern science (cf. *ibid.*). For another, Feigl's is an inclusive criterion of reality sharply different from verificationism's. It states that to assert that something is real is to give it a place within the spatio-temporal-causal framework of science.

Feigl was sensitive to the idea that the adoption of scientific realism – and hence the concomitant criterion of reality – is, *ultimately*, a matter of convention: it is based on a decision to expand the conceptual framework through which we theorize about the world. This decision, he argued, required a *Copernican turn*. Whereas empiricism had traditionally started with the world of experience and had aimed to show how the object of science should be made to fit within the object of perception, realism should take the object of perception to fit within the object of science. Better put, perception is epistemically special, because it is through this that human beings get to know what the world is like, but the data of perception (as well as the perceivers) are part of the natural world, as this is described by science, and the question is how they fit into the thus described natural world. So, the move from empiricism to realism requires a change of perspective, which is not dictated by reason or evidence.

Are, then, electrons and the like *real*? The answer is clearly 'yes' if it is seen as being asked within the framework of realism. But if there were some further anxiety as to whether electrons and the like were *really* real, it would have to be quelled. Feigl shared the view with Carnap that if we take the empiricist critique of traditional metaphysics seriously, there is no framework-free standpoint from which what there is can be viewed. The question of what there is (better: the question of what one is committed to) can only be settled *within* a framework, and its answer has to do with what types of entity have to be assumed for the framework to play the role it is supposed to.

This kind of rapprochement requires the view that metaphysical (or ontological) issues can be clearly and forcefully distinguished from scientific ones. This was precisely Carnap's view. Carnap (1950) argued that ontological questions could be asked in two distinct ways: as external questions and as internal ones. He went on to exclude *external theoretical questions*: questions about the reality of a general *type* (or category) of entity which are supposed to be settled by looking for (empirical) evidence for the reality of this type or by insight into the metaphysical structure of the world. Questions concerning the reality of a type of entity are legitimate and have content, but only if they are taken to be *either* external, *practical* questions concerning the benefits of adopting a certain framework which includes this type of entity in its basis ontic inventory, *or as internal*, theoretical questions concerning the evidence

there is for (or other reasons for accepting the reality of) certain tokens of this type, but only after a framework has been adopted.

Although Quine (1951) was sharply critical of Carnap's distinction, he did agree with Carnap (and Feigl) on a fundamental point, viz., that there is no theory-free standpoint from which what there is can be viewed. For him, however, there is no sharp line between theoretical issues (or questions) and practical ones. Ontological questions (questions about what there is) are theoretical questions *as well as* practical ones: they are answered by our best theory and there is no extra-theoretical court of appeal. The best theory (if indeed there is a unique best theory) just *is* the theory that works sufficiently well – in particular the theory that tallies with the evidence and satisfies a number of virtues, most notably simplicity. For Quine, the utility of a posit and its reality go hand in hand. There is then, no difference between a framework and the theories within it. The framework itself *is* a theory (perhaps a general one) and is judged using the same evidential standards and pragmatic considerations as in the case of ordinary theories. It follows that the entities we are committed to are those that are required for the truth of our overall best theory of the world.

Already in *Two Dogmas of Empiricism*, Quine had argued for the 'blurring of the supposed boundary between speculative metaphysics and natural science'. But this is emphatically not the metaphysics of the traditional metaphysician. Quine takes it that ontological questions are on a par with questions of natural science, but adds that scientific questions are not purely factual, either, in that they, too, concern the choice of a convenient 'scheme or framework for science'. This leads to a blurring of the distinction between the factual and the conventional, and, in turn, it paves the way for a full-blown commitment to the reality of the theoretical entities posited by science. Quine's master argument for the reality of molecules and the like was that they are on a par with the most ordinary physical objects. Hence, the denier of theoretical entities is faced with a *tu quoque*: if you doubt the reality of molecules, you should doubt the reality of the bodies of common sense.

The philosophical issue then becomes the following: given that theories should be taken at face value (which is the gist of what Feigl called 'semantic realism'), can theoretical entities be dispensed with? And besides, given semantic realism, can theories be taken to be true? If the answers are 'no' and 'yes', respectively, (as they were for Feigl, Quine, Sellars and others), the issue of scientific realism is settled. To be sure, the second question cannot be settled in an absolute way. The truth of scientific theories cannot be proved. But the thought was – and it was a great thought – that what matters is the confirmation of scientific theories; given semantic realism, if scientific theories are well-confirmed, there are reasons to believe in the reality of the theoretical entities they posit. As Sellars (1963, p. 97) summed this point up, to have a

good reason for holding a theory is ipso facto to have good reasons for holding that the entities postulated by the theory exist.

4. Ramsey-Sentences: A Failed Truce

A prima facie heavy blow to the ineliminability of theoretical terms came from some unexpected quarters and, in particular, from the application to philosophy of science of what came to be known as *Craig's theorem*: for any scientific theory T, T is replaceable by another (axiomatizable) theory Craig(T), consisting of all and only the theorems of T which are formulated in the observational vocabulary.

This theorem was readily seized upon by instrumentalists of all sorts. Broadly understood, instrumentalism claims that theories should be seen as (useful) instruments for the organization, classification and prediction of observable phenomena. A clear version of this view can be found in Philipp Frank (1932), whose instrumentalism, in modern terminology, is a form of non-cognitivism: theories are symbolic tools that do not (aim to) represent anything which is not antecedently given in experience. One important argument against non-cognitivism is that it is a reconstruction of science that turns a perfectly meaningful practice – where there is communication and understanding – into a meaningless manipulation of symbols underlied by problematic and context-dependent rules that connect some of the symbols with experience (and hence give them some *partial* meaning).

Perhaps for reasons such as this, instrumentalism was taken to require an eliminative dimension, associated with Ernst Mach. The idea is that theoretical discourse is, ultimately, eliminable: whatever in experience can be captured with it, it can be captured without it. This eliminative dimension – which had met only with failures – was given a new breath of life by Craig's theorem. It was argued that theoretical commitments in science were dispensable; theoretical terms could be eliminated *en bloc*, without loss in the deductive connections between the observable consequences of the theory. If so, the question of whether they refer to unobservable entities becomes moot.

This predicament led Hempel (1958) to formulate what he called 'the theoretician's dilemma'. If the theoretical terms and principles of a theory do not serve their purpose of a deductive systematization of the empirical consequences of the theory, they are dispensable. But, given Craig's theorem, even if they do serve their purpose, they can be dispensed with. Hence, the theoretical terms and principles of any theory are dispensable. But is this dilemma compelling? Hempel himself stressed that it is implausible to think of theories as solely establishing a deductive systematization of observable phenomena. He argued that theories also offer inductive systematizations:

they function as premises in *inductive* arguments whose other premises concern observable phenomena and whose conclusion refers to observable phenomena. Accordingly, theories proper – as opposed to their Craig-transforms – become indispensable in establishing inductive connections between observations.

Two opposed reactions to Hempel's dilemma were exemplified in the work of Sellars and Carnap. From the indispensability of theories, Sellars drew the obvious realist conclusion that theoretical entities are real. Challenging the empiricist idea that the major role theories have is to explain empirical generalizations (by deductive or inductive systematization), Sellars offered a direct route to commitment to theoretical entities – via their role in the explanation of singular observable events – and, in particular, via the theory-based explanation of why some observable entities didn't behave the way they should have, had their behaviour been governed by an empirically established observational generalization. In Sellars's view, scientific explanation proceeds via the *theoretical identifications* of observable entities with unobservables. Not only do the latter explain the behaviour of some observable entities; they really *are* the constituents of observable entities. It's not puzzling, then, that if we take our scientific image of the world seriously, we should be committed to unobservables. (For more on this, see Psillos 2004.)

Carnap, on the other hand, resisted till the end the demise of the distinction between science and metaphysics that the Sellarsian move had consolidated. As is now well documented (see Psillos 1999, chapter 3), he found in the Ramsey-sentence approach to theories a way to capture structuralism, and in particular, the thought that the proper content of a theory (beyond its empirical content) is fully captured by the logico-mathematical structure of the Ramsey-sentence of the theory and its existential implications. In this, let's call it 'Ramsey-sentence-structuralism', Carnap thought he found a stone by which to kill two birds. He could defuse the debate between realism and instrumentalism as being merely about a choice of language, and he could secure the proper empirical content of scientific theories against the spectre of metaphysics. The key idea is that the Ramsey-sentence of a theory fares differently from the Craig-transform of the theory; hence it avoids the problems faced by the latter. The Ramsey-sentence^{RT} of a theory T has exactly the same observational content as T; it has exactly the same deductive structure; it can play the same role as T in reasoning by means of theory. Recall, however, that to get the Ramsey-sentence of a theory all predicates which are deemed theoretical are replaced with variables which are bound by an equal number of existential quantifiers. Hence, by its very construction, the Ramsey-sentence dispenses with theoretical predicates; it removes, at least *prima facie*, the issue of the reference of theoretical terms/predicates. Besides, as Carnap was first to note, the very theory T can be written down as a conjunction of two parts: the

Ramsey-sentence ${}^R T$ of T and the conditional ${}^R T \rightarrow T$, which came to be known as *Carnap-sentence*.

The Ramsey sentence ${}^R T$ says that there are classes of entities which are correlated with the observable events in the way the postulates of the theory describe, but it does not say what exactly those entities are; it does not pick out any such class in particular. It can be seen as capturing the structural-cum-empirical content of the theory: what the theory says of the world which can be assessed in terms of truth *and* falsity. Carnap took it to capture the synthetic component of a theory. The Carnap sentence ${}^R T \rightarrow T$ should be read thus: if there are entities that satisfy the Ramsey-sentence, these entities are those that render the theory true. But though the Carnap sentence appears to have genuine empirical content, it does not. Carnap took it to be a meaning postulate, hence, to capture the analytic component of the theory. To be more precise, the Carnap sentence should be seen as a principle constitutive of the conceptual framework of a scientific theory; it defines (implicitly) its theoretical concepts and *ipso facto* the object of knowledge of the theory, viz., whatever satisfies its Ramsey-sentence. (For more on this, see Psillos and Christopoulou 2009).

If all had gone according to plan, Carnap would have achieved two things. First, he would have shown that the difference between Ramsey-sentence-structuralism and realism was only about the Carnap sentence, which is without factual content anyway! Carnap went on to equate Ramsey-sentence-structuralism with instrumentalism and to declare that he had thereby shown that the difference between realism and instrumentalism is essentially linguistic. Second, he would have shown that there would still be room to adopt a form of realism (Ramsey-sentence structuralism) without giving up an essentially anti-metaphysics stance. For the further issue of the supposed excess content of the theory over its Ramsey-sentence (of the *real* reality of theoretical entities, so to speak) turns out to be either an internal, and hence scientifically kosher, issue of what follows from the adoption of a set of implicit definitions offered by the Carnap-sentence, or else an external, and hence metaphysically impotent, issue of choice of a language.

Things didn't go according to plan, however. Ramsey-sentence-structuralism is not the proper way to capture instrumentalism, simply because it yields commitments to the entities (the satisfiers of the Ramsey-sentence) that are not allowed by standard versions of instrumentalism. What is the gain then? Carnap thought there can be a reading of the Ramsey-sentence such that its satisfiers – and hence the commitments that follow from it – are not standard unobservable entities of the sort favoured by realists. He very explicitly took it that, where the Ramsey-sentence says that there are non-empty classes of entities which are related to observable entities by the relations given in the original theory, we are at liberty to think of these classes as classes of *mathematical* entities (cf. 1963, p. 963). His radical view, then, was that Ramsey-sentence-structuralism

was compatible with instrumentalism precisely because it need not imply commitment to the reality of physical unobservable entities. Unless, however, commitment to mathematical entities can be taken to be metaphysically lightweight, the door to metaphysics is wide open again.

In any case, the neutral stance Carnap envisaged faced an unexpected difficulty which came from the fact that, without further assumptions, the truth of the theory collapses to the truth of its Ramsey-sentence. The Ramsey-sentence of a theory can certainly be false, since it might be empirically inadequate. But *if* it is empirically adequate, it cannot be false, provided that the universe of discourse has the right cardinality. Roughly put, if the world has enough objects, (at least as many as required for the truth of the Ramsey-sentence), the variables of the Ramsey-sentence can be assumed to take those objects as values – whatever they are. So, if the Ramsey-sentence is empirically adequate, the only way in which the world might fail to satisfy it is by not having *enough* entities to make the Ramsey-sentence true. The problem is particularly astute for Carnap's Ramsey-sentence neutralism, committed as it was to the satisfiers of the Ramsey-sentence being mathematical entities. It is a priori true that there are always enough of them to satisfy the (empirically adequate) Ramsey-sentence of *any* theory; hence, the very idea of an empirically adequate but false theory becomes an oxymoron – which certainly constitutes a fundamental revision of our conception of theories and of our give-and-take with the world. All this is a version of the much-discussed Newman problem that plagues most versions of structuralism. (For a recent useful discussion, see Ainsworth 2009). The relevant point here is that Carnap's attempt to reconcile empiricism with realism while avoiding metaphysics had to walk the tightrope of Ramsey-sentence structuralism, and its very survival depended very much on what Carnap was willing, in the end, to sacrifice: leaving the image of the world as described by science intact or giving way to metaphysics?

5. Explanation-based Metaphysics

In his review of Jack Smart's groundbreaking work (1963), Quine (1964) exclaimed: 'With science dominating our lives and progressing ever faster on even more frontiers, it is strange that such a view [the realistic view of fundamental particles of physics] needs urging. Strange but true.' By the 1960s, the tide had started to move the realists' way. The agonizing over semantic issues had led to a new consensus: realist (that is, face value) semantics. This makes plausible the claim that theories have 'excess content' over their observational consequences. In light of this, there is a straightforward answer to the following question: what is the world like, according to a given

scientific theory? (Or, equivalently, what is the world like, if a certain scientific theory is true?) The answer is clear and crisp: the world is the way the theory – literally understood – describes it to be.

The move towards explanation-based metaphysics is most clearly seen in Smart's claim that the defence of scientific realism rests on an abductive argument. Smart argued, against instrumentalists, that they must believe in *cosmic coincidence*: a vast number of ontologically disconnected observable phenomena just *happen* to be, and just *happen* to be related to one another, in the way suggested by the theory. Scientific realism, on the other hand, leaves no space for such a coincidence; it is *because* the unobservable entities posited by theories exist that the phenomena are, and are related to one another, the way they are.

This kind of argument pattern bridges the gap between science and metaphysics from the moment it is generally accepted that it is, precisely, inference to the best explanation which is widely used by scientists when they come to accept scientific theories. A recognition such as this is by no means obvious. Pierre Duhem (1906), for instance, put forward an anti-explanationist form of instrumentalism which rested on a sharp distinction between science and metaphysics and claimed that explanation belongs to metaphysics and not to science. Driven by his opposition to atomism and his defence of phenomenological energetics, Duhem envisaged the 'autonomy' of physics, which was seen, by and large, as dependent on a strict conception of the scientific method, captured by the slogan, scientific method=experience + logic. And yet, Duhem went on to offer some of the most powerful arguments in favour of scientific realism, the most central being that the fact that some theories generate *novel* predictions could not be accounted for on a purely instrumentalist understanding of scientific theories. This is a precursor of Smart's argument for realism and, yet, by insisting on the dichotomy between science and metaphysics and by equating metaphysics with the call for explanation, Duhem remained ambivalent as to the status of this argument for realism.

Once it is accepted that this dual stance is deeply problematic precisely because a metaphysics- (that is, explanation-) free science is a chimera, the door is open for the explanationist defence of realism. Hilary Putnam and Richard Boyd argued that, in light of the fact that inference to the best explanation is the very method scientists use to form and justify their beliefs in unobservable entities, scientific realism should be seen as an overarching empirical hypothesis which gets support from the fact that it offers the best explanation of the success of science. The Putnam-Boyd argument came to be known as 'the no-miracles argument' since, in Putnam's (1975, p. 73) slogan,

The positive argument for realism is that it is the only philosophy that does not make the success of science a miracle.

All in all, the realist turn in the philosophy of science made metaphysics legitimate again. But what the empiricist critique of it left behind – despite its overall failure – can be captured by Feigl's nice words: 'if this be metaphysics, make the least of it!' When it comes to the role of metaphysics in the scientific realism debate, it was confined mostly to a declaration of independence (the world that science aims to describe exists in a mind-independent way) and to a commitment to the independent reality of unobservable entities.

This was not an empty gesture, however. It cut a lot of ice against a species of non-sceptical scientific anti-realism which gained some currency in the 1960s and was motivated by the thought that, while the world as described by science should be left intact, this is not necessarily a mind-independent ready-made world. This thought can be traced to the work of the later Wittgenstein and has been advanced by Norwood Russell Hanson (1958). In this view, what there is and what one is committed to depends on the 'logical grammar' of the language one uses to speak of the world, where the 'logical grammar' was meant to capture the interconnections of the uses of key concepts that structure a certain language-game. Science is a 'language-game' which is characterized by its norms, rules, practices and concepts, though all these are *internal* to the game: they don't give the language-users purchase on an independent world. One can then play the science language-game and adhere to its norms and practices. One can follow the scientific method (and in particular the abductive explanatory practices of scientists) and come to accept theories as true as well as believe in the existence of unobservable entities. One, that is, need not be a sceptic. But, on Hanson's view, one need not (perhaps, should not) add to this non-sceptical approach any realist metaphysics. Nor should one build into the language-game a concept of truth that is evidence-transcendent.

The right realist answer to this challenge was to emphasize that the world comes already structured. That the world has a built-in natural structure is licensed as the best explanation of the friction there is between the world and our scientific theories or paradigms. The presence and persistence of anomalies in scientific theories is best explained by the fact that there is a mismatch between the actual natural structure of the world and the ways in which this structure is modelled by theories.

6. The Battle of Empiricism – Phase II

Non-sceptical anti-realism never became too popular among philosophers of science. The main rival of scientific realism in the last decades of the twentieth century was van Fraassen's (1980) constructive empiricism. The debate took a distinctively epistemic turn – though this was occasionally disguised by the

fact that van Fraassen characterized both scientific realism and constructive empiricism in primarily axiological terms: realism takes it that science aims at truth, while constructive empiricism takes science as an activity that aims at empirical adequacy. In the background of this axiological characterization was a full endorsement of realist semantics. This endorsement, to be sure, did not dictate acceptance of scientific theories as true or truthlike; it is consistent for an empiricist to suspend belief in the truth of accepted scientific theories and take it that they can, at best, be assessed in terms of empirical adequacy. Why, however, would someone opt for this view unless one believed that truth was either unachievable or at least unnecessary for science?

There is a kind of oscillation between these two views. If it is claimed that truth is unachievable, some positive reason should be offered for this. In particular, the reason should be such that it challenges the ability of scientific method to produce a well-confirmed account of the unobservable structure of the world, and hence it undermines the rationality of belief in such an account. If, however, it is merely claimed that truth is unnecessary – in the sense that science can be made sense of without being taken to deliver truth (about the unobservable world) – belief in truth ends up neither irrational nor unwarranted. Accordingly, there are two ways to view the recent empiricist attempts to resist scientific realism.

The weak way (*viz.*, the view that looking for truth is unnecessary) is meant to be ecumenical. Realism can coexist with constructive empiricism; neither of them is rationally compelling. But this weak way to resist realism faces two problems. The first is that constructive empiricism can be flanked from the left, as it were. For there are weaker positions available that render supererogatory even belief in the empirical adequacy of a theory. For instance, it could be argued that science aims at unrefuted theories and that acceptance of a theory involves only the belief that it is unrefuted. But an unrefuted theory is not necessarily an empirically adequate theory. The second problem is the one faced by Ramsey-sentence structuralism. As has been noted recently by Demopoulos (2003) and Ketland (2004), even with a model-theoretic understanding of empirical adequacy, like the one adopted by van Fraassen, unless further constraints are imposed on the models that render true an empirically adequate theory, empirical adequacy collapses to truth, in the sense that an empirically adequate theory cannot fail to be true. These extra constraints allow for the possibility that the world might not be among the models that satisfy an empirically adequate theory. If, however, such extra constraints are imposed – if, for instance, it is acknowledged that the world has a certain natural unobservable structure which might not be captured by an otherwise empirically adequate theory – constructive empiricism loses out on two counts: it puts its official agnosticism about the unobservable structure of the world in jeopardy and makes a rather significant metaphysical concession

to realism. A notable irony here is that, despite their eminent differences, the non-sceptical version of anti-realism, Carnap's Ramsey-sentence-structuralism and constructive empiricism, *all* (and for different reasons) have to come to terms with the realist claim that the world has a certain natural structure.

What if constructive empiricism resists realism the hard way? What if it bases its resistance on the strong claim that the truth (about the unobservable) is unachievable or unavailable? This strong way to resist realism would be sectarian. Realism can no longer coexist with constructive empiricism: realism would be rationally bankrupt. But this strong way to resist realism faces two problems. The first is simply that constructive empiricism would also end up being rationally bankrupt. If the problem with realism was that theories cannot be proved to be true, the very same problem would hold for attempts to prove that theories are empirically adequate. Unless it is shown that assertions about unobservables are confirmed, or otherwise tested in ways that are essentially different from the ways in which assertions about observables are confirmed or otherwise tested, constructive empiricism would be no less precarious than realism. The second problem is that all attempts to show that there is a principled epistemic difference between observables and unobservables have been found wanting.

7. Resisting Epistemic Dichotomies

Van Fraassen's critique of scientific realism was premised on making a natural distinction between observable and unobservable entities carry the weight of a sharp epistemic dichotomy between those aspects of nature that are knowable and those that are not. It is extremely interesting that many realists followed suit and developed positions that rested on epistemic dichotomies, which, however, were drawn *within* the realm of the unobservable. The key thought was that there is no problem with having epistemic access to the unobservable in general, but there is a problem with having such access to some aspects of it, or types of it, and so forth. A main reason for this selective scepticism comes from the so-called pessimistic induction. Before we make this link specific, let us have a summary of the most salient epistemically dichotomous positions that are meant to challenge scientific realism.

(Epistemic) Structural realism: The epistemic dichotomy is between knowing the structure of nature and knowing whatever is left to fill out the structure (the unobservable 'fillers' of the structure of the world). This is an epistemic distinction among bits of the unobservable world – its structure and its non-structure.

Entity realism: The epistemic dichotomy is between knowing entities (and perhaps some of their properties) and knowing the truth of (fundamental) theories.

Semi-realism: The epistemic dichotomy is between detection properties of particulars, that is properties of concrete causal structures, and auxiliary properties, that is properties attributed to particulars by theories, but for which there is no reason to believe in their reality, since they are not detected (though they might be detectable and become detected later on; cf. Chakravartty 2007).

Neo-instrumentalism: The epistemic dichotomy is between those entities to which there is an independent route of epistemic access (mediated by theories that cannot be subjected to serious doubt) and those entities to which all supposed epistemic access is mediated by high-level theories (cf. Stanford 2006).

As noted already, all these positions draw the epistemic dichotomy *within* the realm of the unobservable, therefore allowing that there is epistemic access at least to *some* unobservable parts of reality. That some knowledge of the unobservable is deemed possible is an epistemic victory for realism! But might it not be a Pyrrhic victory?

Not quite! The common denominator of all these epistemically dichotomous positions is precisely this: that there is a principled epistemic division between what can be known of nature and what cannot. So, there is a principled limit to the scientific knowledge of the world. The limit is different in the assorted positions, but it is always *principled*, *definite* and drawn by philosophical reflection and argument. What exactly, then, is the *philosophical* issue at stake? I take it to be this: given the ineliminability of theories from science; given the realist reading of the semantics of theories; given that theories, if true, imply commitment to unobservables; are there reasons to accept the existence of a strict, sharp, robust, and principled dichotomy between the epistemically accessible unobservables and the epistemically inaccessible ones? Only science can tell us what the world is like. Philosophy can only raise some principled challenges to the ability of science to tell us what the world is like. Science might, in the end, not succeed in revealing what the world is like. It might be able to disclose only part of the structure and furniture of the world. But this is as it should be. It would be a totally different matter if there were good reasons – mostly drawn by philosophical reflection on science, its methods and its limits – to believe that we, *qua* cognitive beings, or science *qua* an epistemic enterprise, are cognitively closed, in a principled manner, to some aspects of the unobservable world.

But there is no *good* reason (either a priori or a posteriori) to impose a principled epistemic division between what can be known of nature and

what cannot. There might be parts of nature that science might never be able to map out, but these do not fall nicely within a conceptual category which captures one side of a sharp epistemic dichotomy (the unknown X : the unobservable; the non-structure; the intrinsic properties; the auxiliary properties; whatever-there-is-only-theory-mediated-access-to, and the like). The argument for this realist reaction is book-length – dealing, as it has to, with all current attempts to resist realism. It is offered in Psillos (2009).

As noted already, an important motivation for resisting the epistemic optimism associated with scientific realism – the view that science does succeed in offering a truth-like account of the world – has come from the past failures of scientific theories. The key philosophical move that shaped much of the discussion in the last quarter of the twentieth century was based on the claim that there is no reason to think that current scientific theories enjoy any epistemic privilege over their abandoned predecessors. The so-called Pessimistic Induction over the history of science capitalized on the fact that, despite their empirical successes, many past theories were abandoned and replaced by others. Would it, then, not be natural to conclude that the current ones will face that same fate in due course?

Philip Kitcher (1993) and I (1999) have aimed to resist this ‘natural’ conclusion, by arguing that there are ways to distinguish between the ‘good’ and the ‘bad’ parts of past abandoned theories and by showing that the ‘good’ parts – those that enjoyed evidential support, were not idle components and the like – were retained in subsequent theories. This kind of response suggests that there has been enough theoretical continuity in theory-change to warrant the realist claim that science is ‘on the right track’. To be more precise, the realist strategy proceeds in two steps. The first is to make the claim of continuity (or convergence) plausible, *viz.*, to show that there is continuity in theory-change, and that this is not merely empirical continuity: substantive theoretical claims that featured in past theories and played a key role in their successes (especially novel predictions) have been incorporated in subsequent theories and continue to play an important role in making them empirically successful. But this first step does not establish that the convergence is to the truth. For this claim to be made plausible, a second argument is needed, *viz.*, that the emergence of this evolving-but-convergent network of theoretical assertions is best explained by the assumption that it is, by and large, approximately true.

Note that though this kind of realism is selective – in that it does *not* imply belief in everything a theory implies – its selectivity is not the outcome of a philosophically driven adherence to an epistemic dichotomy of the kind noted above. It is one thing to accept the (eminently plausible) view that not all parts of a scientific theory are confirmed by the evidence; it is quite another thing to impose a principled distinction between those parts that are confirmed and those that are not, which is supposed to mirror (or constitute) a sharp epistemic

dichotomy. Worrall's (1989) structural realism – which initiated this selective realist epistemic attitude – should be blamed for a jump, from the sensible view that the actual historical development of theories should teach realists the lesson of being selective in what part of theories they take seriously as warrantably telling us what the world is like, to the controversial conclusion that only the mathematical structure of the theory can warrantably tell us what (the structure of) the world is like.

The emergent consensus is that the unobservable is not, in principle, epistemically inaccessible. Actually, science has succeeded in telling us a lot about a lot of unobservables – a lot that we can warrantably take to be part of a stable and broadly truthlike scientific image of the world. Why should we have expected anything more from the scientific realism debate?

8. From Neo-Humeanism to Neo-Aristotelianism

The story told so far has been premised on the assumption that it is not philosophy's job to revise the description of the world, as this is offered by our best scientific theories, but rather to interpret these theories and tease out what the world is like according to them. The empiricists' critique of metaphysics didn't defy this assumption, though, as we have seen, it didn't invariably succeed in steering a steady course between anti-revisionism and anti-metaphysics. But to some empiricists (notably Duhem and van Fraassen), the critique of metaphysics is tied to the critique of explanation by postulation – that is, explanation in terms of unobservable entities and mechanisms. This is supposed to be the pinnacle of inflationary metaphysics.

Why, one may wonder, is explanation-by-postulation inflationary? In a sense, it obviously is: it proceeds by positing further entities that are meant to *explain* the life-world and its (typically non-strict) laws. But in another sense, it is not. For, if you think of it, it proceeds by positing micro-constituents of macro-objects, whose main difference from them is that they are, typically, unobservable. That a putative entity is unobservable is, if anything, a relational property of this entity and has to do with the presence of observers with certain sensory modalities (of the kind we have), and not others. No interesting metaphysical conclusions follow from this fact, nor any seriously controversial ontological inflation.

If the issue is to avoid metaphysics, a lot depends on how much of it we should want to avoid and how. The empiricist tradition we examined in Sections 2 and 3 didn't take an anti-metaphysical stance to imply (or to require) what Feigl once called the 'phobia of the invisible and the intangible'. This tradition, which needs to be resurrected and further defended, avoids metaphysics not by taking a revisionist stance towards modern science (replete as

it is with talk about unobservables) but by noting that, at the end of the day, fundamental metaphysical questions are framework questions and are not dealt with in the same way in which questions about the reality of ordinary entities (be they stones or electrons) are dealt with – the relevant ontic framework must already be in place before questions about the reality of specific entities are raised. In light of this, an empiricist can accept the realist framework – in essence, a framework that posits entities as constituents of the commonsensical entities and relies on them and their properties for the explanation and prediction of the laws and the properties of commonsensical entities – without metaphysical anxieties. (For more on this, see Psillos forthcoming a). Here again, there is a nice compromise: along with empiricism's anti-metaphysical stance, we accept that there is no framework-free standpoint from which what there is can be viewed; but along with realism, we accept that what there is what is required for a coherent and unified causal-nomological scientific image of the world.

There *is*, to be sure, some residual metaphysical anxiety to be quelled. This is partly due to the following predicament: if explanation is the way to do metaphysics, and if metaphysical commitments, like ordinary scientific ones, are the product of the application of inference to the best explanation, where should we stop? Why not accept all the really metaphysically inflationary armoury of the neo-Aristotelians? Indeed, an increasing number of realists wed realism with a neo-Aristotelian view of the deep structure of reality, based on the claim that commitment to this rich metaphysics, which leaves behind the Humean barren landscapes, is licensed by the very method by means of which scientists form and justify their beliefs in the unobservable, *viz.*, inference to the best explanation (cf. Sankey 2008). Where, exactly, one stops seems to be a matter of taste (and perhaps of some argument as to what exactly *is* the best explanation of whatever the explanandum is taken to be). But many (e.g. Brian Ellis and Alexander Bird) go all the way and adopt metaphysical accounts of causation and laws and dispositional essentialism about properties and natural kinds. As Sankey acknowledges, these are 'optional doctrines' of scientific realism. So why should they be bought? It's not clear to me why. It is one thing to say, for instance, that the world has a pre-existing and determined natural structure – implying as it does that the world is not structurally amorphous. This might well be based on admitting objective similarities and differences in nature – even natural groupings of properties. It is quite another thing to hypostatize kinds or to be an essentialist about them. The same attitude could be had regarding laws of nature or universals or powers and potencies and all the other forbidden fruits of the neo-Aristotelian garden of Eden (for more on this, see Psillos 2005).

One possible problem with this heavyweight neo-Aristotelian conception of scientific realism stems from the fact that some of its advocates have

subscribed to an epistemically dichotomous position, noted above. This creates an important tension – which is very clearly seen in Chakravartty's case. He, like many others, subscribes to the full panoply of neo-Aristotelianism. At the same time, he takes it that scientific realists should be committed only to the detection, as opposed to the auxiliary, properties of particulars. None of the extra stuff that Chakravartty finds in the world (*de re* necessities, ungrounded dispositions and the like) are detectable. They are taken to be part of the baggage of scientific realism because they play a certain explanatory role, notably, they distinguish causal laws from merely accidental regularities. So, we are invited to accept a certain set of double standards – one for scientific theories, and another for metaphysics. While in the case of scientific theories, epistemic optimism requires causal contact with the world, thus denying epistemic optimism merely on the basis of the explanatory virtues of theories, in the case of the metaphysical foundations of scientific realism, epistemic optimism is solely the function of explanatory virtues. To put the point somewhat provocatively, the metaphysics of scientific realism ends up being an auxiliary system whose detection properties are Humean regularities and other metaphysically less fatty stuff.

It might be concluded that if one wants to be a neo-Aristotelian scientific realist, one had better not rest one's epistemic attitude towards theories on a too demanding criterion – and in particular one that cannot be honoured by metaphysical theories. Should, in any case, a realist adopt neo-Aristotelianism simply on the basis that it is the best explanation of, say, the neo-Humean account of the world? My own view on this matter is still developing, but in broad outline it comes to this. If we take IBE seriously, as we should, the answer to the above question should be positive. But, it can be contested that neo-Aristotelianism does indeed meet the best explanation test. One particularly acute problem is that all these denizens of the neo-Aristotelian world (powers, metaphysical necessities, dispositional essences and the like) are themselves unexplained explainers. Though everyone should accept some unexplained explainers, in this particular case, they are more poorly understood than the Humean facts that they are supposed to explain. Another problem is that it is not clear at all how all these heavy metaphysical commitments are related to current scientific theories. They are not born out of current theories. Actually, no particular science, let alone particular scientific theory, can yield interesting metaphysical conclusions, simply because each science has its own specific and particular subject matter, whereas the object of metaphysics (at least as understood by many neo-Aristotelians) is very general and domain-independent: it is the fundamental deep structure (or building blocks) of reality as a whole, abstracting away from its specific scientific descriptions.

Accordingly, neo-Aristotelian scientific realists face a dilemma. They have to proceed top-down, that is, to start from an a priori account of the possible

fundamental structure of reality and then try to mould the actual world as described by the sciences into it. The price here is that there is a danger of neglecting or overlooking important differences between sciences or scientific theories in the ways the world is described and in the commitments they imply. For instance, even though physical kinds might conform to the model of dispositional essentialism, biological kinds might not. Alternatively, they have to proceed bottom-up, that is to start with individual sciences or theories and try to form a unified account of the actual deep structure of reality by generalization or abstraction. The price here is that there is no guarantee that this account can be had.

Be that as it may, this is an important item on the agenda of the current debate over scientific realism. Far from being a dirty word, the m-word is hardly dispensed with. Again, however, at stake is how much of metaphysics we should buy into. In between the neo-Aristotelians and the neo-Humeans are the so-called ontic structuralists. A central motivation for ontic structuralism comes from problems in the interpretation of quantum mechanics and the metaphysics (and the physics) of individuality. But the antidote to issues of metaphysical underdetermination is metaphysical, too; to say (roughly) that objects consist entirely of relationships, or that objects do not exist independently of relational structures, hypostatizes relational structures and makes them the ultimate building blocks of reality. Even critics of metaphysics as a whole, like Newton da Costa and Steven French (2003) and Ladyman and Ross (2007), unashamedly take it that structure is *all there is*. Flirting with neo-Aristotelianism, (as in James Ladyman 2001), ontic structuralism might well create an explosive metaphysical brew, simply because it will have to amalgamate in the same position the thought that structures are abstract entities with the thought that they are nonetheless the locus of modality and causality.

9. *Quo Vadis?*

With more than a century of heated debate behind it, one would have expected that the scientific realism debate should have been exhausted. But it isn't, despite some prematurely issued death certificates. The fact is that the fight over scientific realism is as fierce as ever, though the battleground has shifted to the more murky waters of metaphysics. It remains to be seen whether the resurgence of Aristotelian metaphysics will show that the Humean image of the world, mostly associated with the advent of the scientific revolution, will be seen as a small interruption in an essentially Aristotelian image of the world. I very much hope this won't happen, but its very possibility shows that the scientific realism debate remains, as ever, in contact with the most fundamental philosophical issues.

Note

* Many thanks to my good friend Panagiotis Oulis for the Humean twist in the title.

References

- Ainsworth, P. (2009), 'Newman's objection', *The British Journal for the Philosophy of Science*, 60, 135–71.
- Carnap, R. (1937a), 'Testability and meaning – continued', *Philosophy of Science*, 4, 1–40.
- (1937b), *The Logical Syntax of Language*. London: Routledge and Kegan Paul.
- (1950), 'Empiricism, semantics and ontology', *Revue Internationale de Philosophie*, 4, 20–40.
- (1963), 'Replies and systematic expositions', in P. Schilpp (ed.), *The Philosophy of Rudolf Carnap*. La Salle, IL: Open Court.
- Chakravartty, A. (2007), *A Metaphysics for Scientific Realism – Knowing the Unobservable*. Cambridge: Cambridge University Press.
- da Costa, N. and French, S. (2003), *Science and Partial Truth*. Oxford: Oxford University Press.
- Demopoulos, W. (2003), 'On the rational reconstruction of our theoretical knowledge', *The British Journal for the Philosophy of Science*, 54, 371–403.
- Duhem, P. (1906), *The Aim and Structure of Physical Theory* (2nd edn 1914, P. Wiener trans. 1954). Princeton, NJ: Princeton University Press.
- Feigl, H. (1950), 'Existential hypotheses: realistic versus phenomenalist interpretations', *Philosophy of Science*, 17, 35–62.
- Frank, P. (1932), *The Law of Causality and its Limits*, M. Neurath and R. S. Cohen trans. Dordrecht: Kluwer.
- Hanson, N. R. (1958), *Patterns of Discovery: An Inquiry into the Conceptual Foundations of Science*. Cambridge: Cambridge University Press.
- Hempel, C. (1951), 'The concept of cognitive significance: a reconsideration', *Proceedings and Addresses of the American Academy of Arts and Sciences*, 80, 61–77.
- (1958), 'The theoretician's dilemma: a study in the logic of theory construction', *Minnesota Studies in the Philosophy of Science*, 2. Minneapolis, MN: University of Minnesota Press.
- Ketland, J. (2004), 'Empirical adequacy and ramsification', *The British Journal for the Philosophy of Science*, 55, 287–300.
- Kitcher, P. (1993), *The Advancement of Science*. Oxford: Oxford University Press.
- Ladyman, J. (2001), 'Science, metaphysics and structural realism', *Philosophica*, 67, 57–76.
- Ladyman, J. and Ross, D. (2007), *Every Thing Must Go: Metaphysics Naturalised*. Oxford: Oxford University Press.
- Psillos, S. (1999), *Scientific Realism: How Science Tracks Truth*. London: Routledge.
- (2004), 'Tracking the real: through thick and thin', *The British Journal for the Philosophy of Science*, 55, 393–409.
- (2005), 'Scientific realism and metaphysics', *Ratio*, 18, 385–404.
- (2008), 'Carnap and incommensurability', *Philosophical Inquiry*, 30, 135–56.
- (2009), *Knowing the Structure of Nature*. London: Palgrave Macmillan.

- (forthcoming a), 'Choosing the realist framework', *Synthese*, DOI 10.1007/s11229-009-9606-9.
- (forthcoming b), 'On Reichenbach's argument for scientific realism', *Synthese*, DOI 10.1007/s11229-009-9594-9.
- Psillos, S. and Christopoulou, D. (2009), 'The a priori: between conventions and implicit definitions', in N. Kompa, C. Nimtz and C. Suhm (eds), *The A Priori and Its Role in Philosophy*. Cincinnati, OH: Mentis, pp. 205–20.
- Putnam, H. (1975), *Mathematics, Matter and Method*. Cambridge: Cambridge University Press.
- Quine, W. V. O. (1951), 'Carnap's views on ontology', *Philosophical Studies*, 2, 65–72.
- (1964), 'Science and truth', *New York Review of Books*, 2, no. 11, 9 July 1964.
- Reichenbach, H. (1938), *Experience and Prediction*. Chicago, IL: University of Chicago Press.
- Sankey, H. (2008), *Scientific Realism and the Rationality of Science*. Aldershot: Ashgate.
- Schlick, M. (1932), 'Positivismus und realismus', *Erkenntnis*, 3, 1–31, trans. as 'Positivism and realism', in A. J. Ayer (ed.), *Logical Positivism*, 1960. Glencoe, NY: Free Press.
- Sellars, W. (1963), *Science, Perception and Reality* (reissued 1991), Atascadero, CA: Ridgeview Publishing Company.
- Smart J. J. C. (1963), *Philosophy and Scientific Realism*. London: Routledge and Kegan Paul.
- Stanford, P. K. (2006), *Exceeding Our Grasp: Science, History, and the Problem of Unconceived Alternatives*. Oxford: Oxford University Press.
- van Fraassen, B. C. (1980), *The Scientific Image*. Oxford: Clarendon Press.
- Worrall, J. (1989), 'Structural realism: the best of both worlds?' *Dialectica*, 43, 99–124.