Contents

1
The Present State of the Scientific Realism Debate

> The unique attraction of realism is the nice balance of feasibility and dignity that it offers to our quest of knowledge. (...) We want the mountain to be climbable, but we also want it to be a real mountain, not some sort of reification of aspects of ourselves.
>
> Crispin Wright (1988, 25)

Once upon a time there was a feeling in the philosophy of science community that the scientific realism debate had run out of steam. Arthur Fine went as far as to declare that 'realism is well and truly dead' (1986a, 112) and to compose the obituary of the debate, *aka* the Natural Ontological Attitude. Fortunately, the allegations of premature death failed to persuade many philosophers, for whom the scientific realism debate has had a glorious past and a very promising future. In the last dozen of years only there have been a number of books which cast a fresh eye over the issue of scientific realism, such as those by Suppe (1989), Putnam (1990), Almeder (1992), Wright (1992), Kitcher (1993a), Aronson, Harré & Way (1994), Brown (1994), Laudan (1996), Leplin (1997), Kukla (1998), Trout (1998), Cartwright (1999), Giere (1999), Niiniluoto (1999) and Psillos (1999). Although these books differ vastly in their approaches and in their substantive theses, they can all be seen as participating in a common project: to characterise carefully the main features of the realism debate and to offer new ways of either exploring old arguments or thinking in novel terms about the debate itself.

In this chapter I discuss the present state of the scientific realism debate with an eye to important but hitherto unexplored suggestions and open issues that need further work.

1.1 Modesty and presumptuousness

In Psillos (1999), I offered the following three theses as constitutive of scientific realism. Each of these is meant to warn off a certain non-realist approach.

The Metaphysical Thesis: The world has a definite and mind-independent structure.

The Semantic Thesis: Scientific theories should be taken at face value. They are truth-conditioned descriptions of their intended domain, both observable and unobservable. Hence, they are capable of being true or false. The theoretical terms featuring in theories have putative factual reference. So, if scientific theories are true, the unobservable entities they posit populate the world.

The Epistemic Thesis: Mature and predictively successful scientific theories *are* well confirmed and approximately true of the world. So, the entities posited by them, or, at any rate, entities very similar to those posited, inhabit the world.

The *first* thesis means to make scientific realism distinct from all those anti-realist accounts of science, be they traditional idealist and phenomenalist or the more modern verificationist accounts of Michael Dummett's (1982), and Hilary Putnam's (1981, 1990) which, based on an epistemic understanding of the concept of truth, allow no divergence between what there is in the world and what is licensed as existing by a suitable set of epistemic practices, norms and conditions. It implies that if the unobservable natural kinds posited by theories exist at all, they exist independently of the scientists' ability to be in a position to know, verify, recognise, and the cognate, that they do.

The *second* thesis – semantic realism – renders scientific realism different from *eliminative instrumentalist* and *reductive empiricist* accounts. Eliminative instrumentalism (most notably in the form associated with Craig's Theorem) takes the 'cash value' of scientific theories to be fully captured by what theories say about the observable world. This position typically treats theoretical claims as syntactic-mathematical constructs which lack truth-conditions, and hence any assertoric content. Reductive empiricism treats theoretical discourse as being disguised talk about observables and their actual (and possible) behaviour. It is consistent with the claim that theoretical assertions have truth-values, but understands their truth-conditions *reductively*: they are fully captured in an observational vocabulary. Opposing these two positions, scientific realism is an 'ontologically inflationary' view. Understood realistically, the theory admits of a literal interpretation, namely, an interpretation

in which the world is (or, at least, can be) populated by a host of unobservable entities and processes.

The *third* thesis – epistemic optimism – is meant to distinguish scientific realism from *agnostic* or *sceptical* versions of empiricism (cf. van Fraassen 1980, 1985). Its thrust is that science *can* and *does* deliver theoretical truth[1] no less than it can and does deliver observational truth. It's an implicit part of the realist thesis that the ampliative–abductive methods employed by scientists to arrive at their theoretical beliefs are reliable: they tend to generate approximately true beliefs and theories.

Semantic realism is not contested any more. Theoretical discourse is taken to be irreducible and assertoric (contentful) by all sides of the debate. Making semantic realism the object of philosophical consensus was by no means an easy feat, since it involved two highly non-trivial philosophical moves: *first*, the liberalisation of empiricism with its concomitant admission that theoretical discourse has 'excess content', that is, content that cannot be fully captured by means of paraphrase into observational discourse; and *second*, a battery of indispensability arguments which suggested that theoretical terms are indispensable for any attempt to arrive, in Rudolf Carnap's (1939, 64) words, at 'a powerful and efficacious system of laws' and to establish an inductive systematisation of empirical laws (cf. Hempel 1958).

Given this, the distinctive of scientific realism is that it makes two claims in tandem, one of which (to explore Wright's (1992, 1–2) terminology) is 'modest', while the other is more 'presumptuous'. The *modest* claim is that there is an independent and largely unobservable-by-means-of-the-senses world that science tries to map. The *presumptuous* claim is that although this world is independent of human cognitive activity, science succeeds in arriving at a more or less faithful representation of it, that is of knowing the truth (or at least *some* truth) about it.

For many philosophers, this is *ab initio* an impossible combination of views: if the world is independent of our abilities or capacities to investigate it and to recognise the truth of our theories of it, how can it possibly be knowable? Two options then appear to be open to prospective realists: either to compromise presumptuousness or to compromise modesty.

1.1.1 Compromising presumptuousness

Here the cue is taken from Karl Popper's (1982). Take realism to be a thesis about the aim of science (truth), leaving entirely open the issue

of whether this aim is (or can ever be) achieved. Implicit in this strand is that truth is understood realist-style (in the sense of correspondence with the world) in order not to compromise modesty as well. Popper made the headlines by claiming that scientists can never say that this aim has been achieved, but that truth is somehow magically approached by the increasing verisimilitude of successive theories; *magically* because there is nothing in Popper's account of verisimilitude which, even if it worked,[2] guarantees that there is a 'march on truth'. As we shall see in Chapter 3, Musgrave (1996, 23) agrees with Popper that realism is primarily an *axiological thesis*: science aims for true theories. There is clear motivation for this compromise: even if all theories scientists ever come up with are false, realism isn't threatened. Musgrave doesn't think that all scientific theories have been or will be outright false. But he does take this issue (whatever its outcome may be) to have no bearing on whether realism is a correct attitude to science.[3] There are, however, inevitable philosophical worries about the axiological characterisation to realism. First, it seems rather vacuous. Realism is rendered immune to any serious criticism that stems from the empirical claim that science has a poor record in truth-tracking (cf. Laudan 1984). Second, aiming at a goal (truth) whose achievability by the scientific method is left unspecified makes its supposed regulative role totally mysterious. Finally, all the excitement of the realist claim that science engages in a cognitive activity that pushes back the frontiers of ignorance and error is lost.

It seems irresistible that the only real option available for presumptuous realists is to compromise their modesty: if the world isn't in any heavyweight way independent of us, its knowability can be safeguarded. Compromising modesty is typically effected by coupling realism with an epistemic notion of truth which *guarantees* that the truth does not lie outside our cognitive scope.

1.1.2 Compromising modesty

Here the main cue is taken from Putnam's (1981). Take realism to involve an epistemic conception of truth, that is, a conception of truth which guarantees that there cannot be a divergence between what an ideal science will assert of the world and what happens (or there is) in the world. This line has been exploited by Brian Ellis (1985) and Nicholas Jardine (1986). For Ellis, truth is 'what we should believe, if our knowledge were perfected, if it were based on total evidence, was internally coherent and was theoretically integrated in the best possible way' (1985, 68). There are many problems with this view that I won't

rehearse here (cf. Newton-Smith 1989b; Psillos 1999, 253–5). The only thing to note is that it's not obvious at all whether the suggested theory of truth is stable. To use Jardine's (1986, 35) words, the needed concept of truth should be neither too 'secular', nor too 'theological'. It should avoid an awkward dependence of truth on the vagaries of our evolving epistemic values, but it should link truth to *some* notion of not-too-inaccessible epistemic justification. In the attempt to break away from 'secular' notions of truth and to make truth a standing and stable property, it moves towards a 'theological' notion: the justification procedures become so ideal that they lose any intended connection with humanly realisable conditions. In the end, the required epistemic conception of truth becomes either 'secular', resulting is an implausible relativism, or 'theological', and hence not so radically different from a (realist) non-epistemic understanding of truth, according to which truth *outruns* the possibility of (even ideal-limit) justification. To be sure, Putnam (1990, viii) has dissociated his views on truth from the (Peircean) ideal-limit theory on the grounds that the latter is 'fantastic (or utopian)'. Still, his proposed alternative which ties ascriptions of truth with the exemplification of 'sufficiently good epistemic situations' fares no better than the Peircean theory vis-à-vis the secular/theological test above. One can always ask: what other than the realist-style truth of a proposition can *guarantee* that the sufficiently good conditions of justification obtain?[4]

There is an interesting dual thesis advanced by Crispin Wright that (a) a suitable epistemic concept of truth does *not* necessarily compromise the modesty of scientific realism and (b) the best hope for the presumptuousness of scientific realism rests on a broadly verificationist (epistemic) understanding of truth. For Wright, scientific realism stands mainly for (a) anti-reductionism and (b) the claim that theoretical discourse is apt for 'representation or fit with objective worldly states of affairs' (1992, 159). The first part of his thesis stems from the thought that the anti-reductionist stance of semantic realism, which treats theoretical discourse as apt for representation, is consistent with a (suitably) 'evidentially constrained' account of truth. This is so because, he claims, scientific realists may accept *both* that theoretical assertions faithfully represent worldly states-of-affairs *and* that these states-of-affairs are 'in principle' detectable (and hence, in principle verifiable). In particular, these worldly states-of-affairs need not be representable in a humanly intelligible way. On this view, the world ends up being independent of human representation (as scientific realism requires), yet the world is in principle detectable, which means that the relevant notion of truth is suitably 'evidentially constrained', and

hence epistemic. (The motto for Wright's verificationist scientific realists would be: there is no in principle undetectable truth.) The second part of Wright's thesis stems from the thought that the realists' epistemic optimism requires that 'the harvest of best methods is (likely to be) truth and may, *qua* so harvested, be reasonably so regarded (1986, 262). But, he goes on, if truth is not taken to be what is 'essentially certifiable by best method' (as a verificationist realist would have it), there is no guarantee that truth is achievable. So, Wright concludes, either the door is left open to a van Fraassen-type sceptic, or to a Quinean pragmatist who 'cashes out' talk of truth is terms of talk about a(n) (everlasting) set of simplicity-guided adjustments in our evolving network of beliefs in response to empirical anomalies.[5]

Wright (1992) has presented a 'minimalist' concept of truth (not to be confused with Paul Horwich's (1990) account) which is characterised by some 'syntactic and semantic platitudes' (e.g., Tarski's T-schema, good behaviour with respect to negation, a 'thin' correspondence intuition, stability and others). Satisfaction of these platitudes (on Wright's proposal) guarantees that a certain discourse with a truth-predicate in it is assertoric (apt for truth and falsity), but leaves open the question whether the concept of truth has a more robust substance. Some realists believe that the concept of truth does have this more robust (non-epistemic) substance which is captured by a 'thick' notion of correspondence with reality, namely, that the *source* of the truth of theoretical assertions is worldly states-of-affair. This notion is taken by realists to be epistemically unconstrained. Wright juxtaposes to this realist notion of truth an epistemically constrained one: 'superassertibility' (1992, 48). He takes it to be the case that superassertibility meets the minimalist requirements noted above, and then asks whether there are features of a discourse which dictate that this discourse needs or implicates a concept of truth *stronger than* superassertibility. He proposes four criteria for judging whether a discourse implicates an epistemic or a non-epistemic conception of truth (over and above the minimalist common ground): extensional divergence, convergence of opinion (or Cognitive Command), the Euthyphro Contrast, and the width of cosmological role. Put in a nutshell, Wright's claim is the following. It may be that truth (realist-style) and superassertibility are extensionally divergent notions (there are truths which are not superassertible and/or conversely). It may be that truth (realist-style) features in best explanations of why there is convergence-of-opinion in a discourse. It may be that the direction of dependence between truth (realist-style) and superassertibility is one-way only: it's because certain statements are true (realist-style) that they

are superassertible and not conversely. And it may be that the statements in a discourse play a wide cosmological role in that their truth (realist-style) contributes to the explanation of assertions and attitudes in other spheres or discourses. This (extremely compact presentation of Wright's seminal idea) leads me to the following conjecture. Even if Wright is right in pointing out that, *prima facie*, scientific realists need not compromise their modesty by adopting an epistemically constrained conception of truth, the very features of the truth-predicate implicated in the assertoric theoretical discourse in science are such that it satisfies all criteria that Wright himself has suggested as pointing towards the operation (or implication) of a (realist-style) concept of truth in a discourse. If this conjecture is right, then the realist aspiration to modesty *ipso facto* implicates a substantive non-epistemic conception of truth.

What about the second part of Wright's thesis, namely, that scientific realists had better adopt an epistemic conception of truth if they are to retain their epistemic optimism? The problem with this suggestion (which Wright recognises and tries to meet) is that a verificationist version of scientific realism brings with it all of the problems that discredited verificationism as a philosophical theory of meaning (and truth). In particular, the viability of Wright's second thesis depends on two premises: first, that radical underdetermination of theories by evidence is a priori impossible; second, that we can make sense of an observation language which is theory-free and which is used to 'cash out' the suitable notion of verifiability. As for the first premise, it seems obvious that the very logical possibility of two or more mutually incompatible theories being empirically equivalent entails (on the assumption that only one of them can be true) that truth doesn't necessarily lie within our cognitive capacities and practices. As for the second premise, if observation is theory-loaded in such a way that we cannot segregate a set of theory-neutral 'observation reports', we cannot even begin to formulate the thesis that theoretical assertions are true in the sense that they are fully verifiable by means of 'observation reports'.

Some realists (e.g., Michael Devitt and Horwich) take scientific realism to be an ontological doctrine which asserts the existence of unobservable entities, and claim that no doctrine of truth is constitutive of realism. Here company is parted, however. Devitt (1984, Chapter 4) argues that insofar as a concept of truth is involved in the *defence* of realism, it should be a correspondence account in order to safeguard that the world is independent in its existence and nature from what we believe. Horwich (1997), on the other hand, after declaring that the scientific

realism debate is about the independence and accessibility of facts about unobservable entities, takes the view that a 'deflationary' conception of truth (which is itself lightweight and metaphysically neutral) is all that is needed for the defence of scientific realism. His core thought is that the truth-predicate doesn't stand for any complex property, but is a quasi-logical device for forming generalisations over propositions.

One can of course pit Devitt's defence of correspondence truth against Horwich's deflationism. But the serious philosophical issue that remains is Horwich's (1997) thesis that the scientific realism debate can be fully stated and explained without any substantive (i.e., non deflationary) concept of truth. In particular, Horwich claims that even when the concept of truth is explicitly mentioned in a realist (or anti-realist) thesis, for example, when realists say that science achieves theoretical truth, or when instrumentalists say that theoretical hypotheses are truth-valueless, or when verificationists say that all truths are verifiable, even *then* it can be captured by a deflationist understanding of truth. But I doubt that this is so easily established. When realists say, for instance, that theoretical discourse should be understood literally, they imply that theoretical assertions shouldn't be taken to be translatable into a vocabulary which commits only to observable states-of-affair. The notion of translatability (or its lack) may inevitably involve reference to sameness (difference) of truth-conditions, which, arguably, are not part of the resources available to the deflationist (cf. Field 1992, 324–5).[6]

1.1.3 Conceptual independence and epistemic luck

Despite attempts to force a compromise on scientific realists, neither modesty nor presumptuousness has to go. From the claim of independence of the world from human cognitive activity it does not follow either that human inquirers are cognitively closed to this world or that when they come to know it, they somehow *constitute* it as the object of their investigation. All depends on how exactly we understand the realist claim of mind-independence. As will be explained in Chapter 2, Section 2.4, it should be taken to assert the logical-conceptual independence of the world: there is no conceptual or logical link between the truth of a statement and our ability to recognise it, assert it, superassert it and the like. The entities that science studies and finds truths about are deemed independent of us (or of mind in general) not in any causal sense, but only in a logical sense: they are not the outcome of, nor are they constituted by (whatever that means), our conceptualisations and theorising. This kind of independence is consistent with the claim that

science and its methodology are causally dependent on the world. In fact, the realists' claim that the scientists' methods of interaction with the world are such that, at least in favourable circumstances, can lead to the formation of warranted beliefs about the 'deep structure' of the world presupposes *causal contact* with the world.

Despite several pages of philosophical argumentation that this contact with the independent world is impossible because it would amount to 'getting out of our skin' (cf. Rorty 1991, 46ff), or because it's 'conceptually contaminated' (Fine 1986b, 151), it's a simple truth that our (inevitably) conceptual give-and-take with the world need not lead to the neo-idealist (or neo-Kantian) thought that the causal structure of the world is a reflection (or projection) of our concepts and intuitions. The independence of the world needn't be compromised. And it cannot be compromised unless one adopts the implausible view that the worldly entities are *causally constituted* as entities by our conceptual and epistemic conditions, capacities and practices. To be sure, realists need to grant that their 'epistemic optimism' that science has succeeded in tracking truth requires a *epistemic luck*: it's not a priori true that science has to be successful in truth-tracking. If science does succeed in truth-tracking, this is a *contingent fact* about the way the world is and the way scientific method and theories have managed to 'latch onto' it (cf. Boyd 1981). Accordingly, the presumptuousness of realism is a contingent thesis that needs to (and can) be supported and explained by argument which shows that the ampliative–abductive methods of science can produce theoretical truths about the world and deliver theoretical knowledge.

If neither modesty nor presumptuousness need compromising, isn't there still an issue as to how presumptuous scientific realism should be? I think we should reflect a bit on what exactly the *philosophical* problem is. I take it to be the following: is there any good reason to believe that science cannot achieve theoretical truth? That is, is there any good reason to believe that after we have understood the theoretical statements of scientific theories as expressing genuine propositions, we can never be in a warranted position to claim that they are true (or at least, more likely to be true than false)? There are some subtle issues here (to which we shall return below and in Chapter 2), but once we view the problem as just suggested, it transpires that what realism should imply by its presumptuousness is that theoretical truth is achievable (and knowable) no less than observational truth.

This last claim may have a *thin* and a *thick* version. The thin version is defended by Jarrett Leplin (1997). His 'Minimal Epistemic Realism' is

the thesis that 'there are possible empirical conditions that would warrant attributing some measure of truth to theories – not merely to their observable consequences, but to theories themselves' (op. cit., 102). As Leplin is aware, many realists would go for a thicker version. This version should imply (and be engaged in the defence of the claim) that the ampliative–abductive methods of science are reliable and do confer justification on theoretical assertions. This thick version is the standing result of Richard Boyd's contribution to the defence of realism. But why do we need it? A 'thin' account cannot issue in rational or warranted belief in the unobservable entities posited by science (and the assertions made about them). All the thin claim asserts is a subjunctive connection between some possible empirical conditions and the truth of some theoretical assertions. This cannot be the litmus test for scientific realism because, suitably understood, it's universally acknowledged as possible. Not only are we given *no* guarantee that this possible connection may be actual (a condition required for the belief in the truth of a theoretical assertion). More importantly, any attempt to offer such a guarantee would have to go via a defence of the method that connects some empirical condition with the truth of a theoretical assertion. Hence, the defence of the rationality and reliability of these methods cannot be eschewed. To me, all this means that the presumptuous strand in the realist thought should be thick. One issue that needs to be explored – as hinted at by Fred Suppe (1989, 340–6) and developed by Kitcher (1993a, Chapter 3) – is how standard epistemological theories of justification, reliability and belief formation can be employed in the realism debate. It may turn out, as I (Psillos 1999, 83–6) and Suppe (1989, 352) believe it does, that the debate on scientific realism is best conducted in the context of broader epistemological theories about the nature of knowledge, justification etc. (This is issue is explored in some detail in Chapter 10.)

So far, we have resisted the claim that the concept of truth involved in scientific realism should be something less than a 'correspondence with reality'. The well-known pressures have led some realists to back down (e.g., Giere 1999, 6). Others, however, have tried to explicate the notion of correspondence in such a way as to remove from it any sense in which it is 'metaphysically mysterious'. Of these attempts, Kitcher's (1993b, 167–9) stands out because he shows that this notion (a) need not commit us to an implausible view that we should (or need to) compare our assertions with the world and (b) arises out of the idea that a fit between representations and reality *explains* patterns of successful action and intervention. A correspondence account of truth is just

another theory which can be judged for (and accepted on the basis of) its explanatory merits.

1.2 Epistemic optimism

It's hard to exaggerate the role that Sellars played in the realist turn during the 1960s. His attack on the 'myth of the given' and his endorsement of the 'scientific image', according to which what is real is what successful scientific theories posit, prioritised scientific theories over folk theories of the 'manifest image' as our guide to what there is (cf. Churchland 1979). (More details about Sellars's arguments for scientific realism are given in Chapter 5.)

Jack Smart (1963, 39) and Maxwell (1962, 18) followed suit by offering arguments for realism based on the explanation of the success of science. If all these unobservable entities don't exist, if theoretical assertions are not well-confirmed and true descriptions of an unobservable world, it isn't possible to explain the empirical success of science and the predicted observed correlations among observable entities. Putnam (1975, 73) turned all this into a famous slogan: realism 'is the only philosophy of science that does not make the success of science a miracle'. Hence, the well-known 'no-miracles' argument for realism (NMA). To be sure, the central thought in this argument is that the realist assertions offer not the only but the *best* explanation of the success of science. Be that as it may, the point of NMA is that the success of scientific theories lends credence to the following two theses: (a) scientific theories should be interpreted realistically and (b) these theories, so interpreted, are well confirmed because they entail well-confirmed predictions. (A relatively detailed discussion of NMA is given in Chapter 3.)

The original authors of NMA didn't place emphasis on *novel* predictions, which are the litmus test for the ability of alternative approaches to science to explain the success of science. For only on a realist understanding, novel predictions about the phenomena *come as no surprise*. Yet, there has been notorious disagreement as to how exactly the conditions of novelty should be understood. A novel prediction has been taken to be the prediction of a phenomenon whose existence is ascertained only after a theory has predicted it. This, however, cannot be the whole story since theories get support also from explaining already known phenomena. It's been suggested (e.g., Worrall 1985) that the 'temporal view' of novelty should be replaced by a 'novelty-in-use' view: a prediction of an already known phenomenon can be use-novel

with respect to some theory provided that information about this phenomenon was not *used* in the construction of the theory. Yet, it's been very difficult to make precise the intuitive idea of 'use novelty'. A fresh analysis comes from Leplin (1997, 77) who explicates 'novelty' by reference to two conditions: 'independence' and 'uniqueness'. The thrust is that a prediction of a phenomenon O is novel for a theory T if no information about O is necessary for the prediction of O by T, and if at the time T explains and predicts O, no other theory 'provides any viable reason to expect' O. If these requirements are satisfied, it's hard to see what other than the relevant truth of the theory T could best explain the novel predictions.[7]

Why has the realist turn come under so much pressure? The main target of the non-realist onslaught has been realism's epistemic optimism. Note that the original Smart–Maxwell formulation of the 'no miracle' argument rested on the assumption that once semantic realism is established, belief in the truth of genuinely successful scientific theories is (almost inevitably) rationally compelling. Van Fraassen's (1980) reaction to this was that the ampliative–abductive methodology of science fails to connect robustly empirical success and truth: two or more mutually incompatible theories can nonetheless be empirically congruent and hence equally empirically successful. Given that at most one of them can be true, semantic realism can still stand but be accompanied by a sceptical attitude towards the truth of scientific theories.

Realists face a dilemma. As W.H. Newton-Smith (1978, 88) pointed out, realists can cling on an 'Ignorance Response' or an 'Arrogance Response'. On the first horn, realists choose to hang on to a realist metaphysics of an independent world, but sacrifice their epistemic optimism. On the second horn, they try to secure some epistemic optimism, but sacrifice the independence of the world by endorsing a view which denies that there are 'inaccessible facts' which can make one of the many rival theories true. In a way, van Fraassen's own attitude amounts to the 'Ignorance Response'.[8] As for the 'Arrogance Response', it's hard to see how one can be a realist and still endorse it. Trimming down the content of the world so that it contains no inaccessible facts leaves three options available (all of which should be repugnant to realists). The first is to re-interpret the empirically equivalent theories so that they are not understood literally and the apparent conflict among them doesn't even arise (an option entertained by some Logical Empiricists). The second (as already noted in Section 1.2) is to adopt an epistemic notion of truth which makes it the case that only one of the empirically equivalent theories passes the truth-test (cf. Jardine

1986). And the third is to argue that all these theories are true, thereby relativising the concept of truth (cf. some time-slice of Quine 1975, 327–8).

Can realists eschew the 'Ignorance Response'? The gist of van Fraassen's challenge is that the explanatory virtues which are part and parcel of the ampliative–abductive methodology of science need not (and perhaps cannot) be taken to be truth-tropic. Hence, any realist hope to forgo the 'Ignorance Response' by grounding their epistemic optimism on explanatory considerations seems to vanish. Not so fast, though.

Putnam's standing contribution to the realist cause is his thought that the defence of realism cannot be a piece of a priori epistemology, but rather part and parcel of an empirical-naturalistic programme which claims that realism is the best empirical hypothesis of the success of science. Capitalising on this thought, Boyd (1981, 1984) embarked on an attempt to establish the accessibility of (and rational belief in) theoretical truth by trying to defend the reliability of ampliative–abductive inferences. This well-known *abductive defence of realism* starts from the fact that the heavily theory-laden scientific methodology is instrumentally reliable (i.e., it yields correct predictions and is empirically successful) and argues that the best explanation of this instrumental reliability is that the background theories (which inform and dictate the methods used by scientists) are approximately true. This is a philosophical (second order) inference to the best explanation (IBE) which suggests that there is a contingent (a posteriori) link between ampliative–abductive methodology (and the concomitant notion of 'best explanation') and truth. It is this argument that grounds the realists' epistemic optimism. It also removes the sting from the rival argument from the underdetermination of theories by evidence (UTE). For two empirically equivalent theories may not be (as a matter of contingent fact) equally good in their explanatory virtues. Hence one of them may well be the best explanation of the evidence and command rational belief.

In any case, UTE rests on two questionable premises: (I) For any theory T there is at least another one incompatible theory T′ which is empirically congruent with T. (II) If two theories are empirically equivalent, they are epistemically equivalent too (i.e., equally confirmed or supported by the evidence). Both premises have been forcefully challenged by realists. Some (e.g., Newton-Smith 1987) have challenged (I) on the grounds that the thesis it encapsulates is not proven. Note, in passing, that realism should be happy with local scepticism. It may turn out that some domains of inquiry (e.g., the deep structure of space-time)

are beyond our ken. Others (e.g., Glymour 1980; Boyd 1981; Laudan & Leplin 1991; Laudan 1996; Psillos 1999) have objected to (II). Here there are, on the face of it, two strategies available. One (IIa) is to argue that even if we take only empirical evidence in the strictest sense as bearing on the epistemic support of the theory, it does not follow that the class of the observational consequences of the theory is co-extensional with the class of empirical facts that can lend support to the theory. An obvious counter-example to the claim of co-extensionality is that a theory can get indirect support by evidence it does not directly entail. The other strategy (IIb) is to note that theoretical virtues are epistemic in character and hence can bear on the support of the theory. Here again there are two options available to realists: (IIb.1) to argue that some theoretical virtues are constitutive marks of truth (e.g., McMullin 1987); or (IIb.2) to argue for a broad conception of evidence which takes the theoretical virtues to be broadly empirical and contingent marks of truth (cf. Boyd 1981; Churchland 1985; Lycan 1988). (IIb.2) is an attractive strategy for two reasons: (a) it challenges the strict empiricist conception of evidence and its relation to rational belief; (b) it removes the apparent tension between modesty and presumptuousness, without also forging an a priori link between theoretical virtues and truth. (IIb.2) is perhaps the most difficult position to defend, but on closer inspection it may well turn out that (IIa) and (IIb.2) are, at root, the very same strategy (cf. Psillos 1999, 171–6).[9]

Not all defenders of realism take the abductive defence of IBE to be central in the defence of realism. There are a few specific problems here and one more general philosophical. The specific problems regard the notion of explanation and the conditions under which it can be balled 'best'. Some realists countenance specific forms of causal explanation (e.g., Salmon (1984) for the so-called common cause principle, or Cartwright (1983) for 'inference to the most probable cause') but deny that they can suitably generalise to engender a blanket notion of IBE. Others (e.g., Lipton 1991) try to provide (descriptively) an account of when a (potential) explanation is best and then to tell a story as to when this explanation licences inference. In the same boat, Ilkka Niiniluoto (1999, 185–92) sketches a formal model of IBE in which the 'best explanation' is linked to the 'best confirmed' hypothesis, given the evidence. Finally, there are those (e.g., Miller 1987) who argue that there cannot be a general mode of inference called IBE, but instead that local ampliative inferences in science are licensed only when they are backed up by 'topic-specific truisms', that is principles which are so entrenched that no-one in the specific domain can seriously deny them. This last

position, however, is sensitive to the issue of what renders these principles 'truisms' if not the fact that they have been arrived at by a legitimate application of IBE. (For more on these matters see Chapter 10. My current views on the structure of IBE are sketched in Psillos 2007a).

What I called the general philosophical problem of the abductive defence of realism has caused a heated discussion. It has been argued (cf. Laudan 1984, 134; van Fraassen 1985, 255; Fine 1986a,b) that the realists' *use* of (a second-order) IBE in defence of realism is circular and question-begging. For, the thought is, it takes for granted the reliability of a mode of inference which is doubted by non-realists. This challenge has led some realists to question the viability of the abductive strategy. Newton-Smith (1989a, 179), for instance, called the realism associated with this strategy 'faded'. And Rom Harré (1988) left behind 'truth realism' and its 'deeply flawed' abductive defence in favour of a methodological strategy, which he called 'policy realism' (cf. also Hendry 1995).

This issue is a focal point of the debate. A proper appreciation of what is at stake presupposes a better understanding of the broader epistemological agendas of the participants. As is explained in detail in Psillos (1999, Chapter 4), the abductive defence of realism proceeds within a broad naturalistic framework in which the charge of circularity loses its bite because what is sought is not justification of inferential methods and practices (at least in the neo-Cartesian internalist sense) but their explanation and defence (in the epistemological externalist sense). It's not as if the abductive defence of realism should persuade a committed opponent of realism to change sides. Strict empiricists, for instance, are not likely to be moved by any defence of IBE, be it circular or straight, precisely because as Ernan McMullin (1994, 100) has noted, they simply choose to tolerate unexplained regularities and phenomena. (One such regularity is that science has been instrumentally reliable and successful.) Van Fraassen's insistence that the explanatory virtues are merely pragmatic is just a further twist to this tolerance to the unexplained. So, strict empiricists deny the abductive defence of realism not so much because it's circular (they would deny a defence of IBE even if it was straight), but mainly because they refrain from accepting the existence of unobservable entities on any grounds that transcend what can be derived from naked-eye observations. But unless this attitude is itself the most reasonable to adopt (something that I doubt), it doesn't follow that IBE is unreasonable.[10] Nor does it follow that the employment of IBE in an abductive defence of the reliability of IBE is question-begging

and unpersuasive. Many (if not all) use modus ponens unreflectively as a sound inferential rule and yet an establishment of the soundness of modus ponens at the meta-level by an argument which effectively uses modus ponens can still explain to them why and in virtue of what features deductive reasoning is sound. In any case, realists vary in the extent to which they adopt the abductive defence of the reliability of IBE. There are those brazen realists, like Boyd, Trout (1998) and myself (Psillos 1999) who take the charge of circularity seriously and try to meet it within a naturalistic perspective. One central thought in this camp is that there is just abduction as the general mode of ampliative reasoning and if this fails, then no ampliative reasoning (and hence no learning from experience) is possible (for more details on this, see Chapter 10). There are the temperate realists (cf. Leplin 1997, 116) who capitalise on the thought that abduction and induction are distinct modes of reasoning and try to argue that IBE is no worse than ordinary inductions which are OK for non-realists. Finally, there are realists (like Brown 1994, Chapter 1) who side-step the charge of circularity and argue that the explanatory story told by realism is just more adventurous and enlightening than alternative stories.

Yet, there is a deep empirical challenge to realism and its abductive defence: the Pessimistic Induction. As Larry Laudan (1984) has pointed out, the history of science is replete with theories that were once considered to be empirically successful and fruitful, but which turned out to be false and were abandoned. If the history of science is the wasteland of aborted 'best theoretical explanations' of the evidence, then it might well be that current best explanatory theories might take the route to this wasteland in due course. Not all realists find this argument threatening. Some (e.g., Devitt 1984) find it simply (and correctly) overstated. Others (e.g., Almeder 1992) take a 'blind realist' stance: at any given stage of inquiry some of our theoretical beliefs are true, yet we can never tell which are such because 'we have no reliable way of determining which of our currently completely authorised beliefs will suffer truth-value revision in the future' (Almeder 1992, 178).[11] What about those of us who think that we should take seriously the Pessimistic Induction and try to meet it?

Although other strategies may be available, the best defence of realism is to try to reconcile the historical record with some form of realism. In order to do this, realists should be more selective in what they are realists about. A claim that now emerges with some force is that theory-change is not as radical and discontinuous as the opponents of scientific realism have suggested. Realists have showed that there are

ways to identify the theoretical constituents of abandoned scientific theories which essentially contributed to their successes, to separate them from others that were 'idle' – or as Kitcher has put it, merely 'presuppositional posits' – and to demonstrate that the components that made essential contributions to the theory's empirical success were those retained in subsequent theories of the same domain (cf. Kitcher 1993a; Psillos 1999, Chapter 5). Given this, the fact that our current best theories may be replaced by others does not, necessarily, undermine scientific realism. All it shows is that (a) we cannot get at the truth all at once; and (b) our judgements from empirical support to approximate truth should be more refined and cautious in that they should only commit us to the theoretical constituents that do enjoy evidential support and contribute to the empirical successes of the theory. Realists ground their epistemic optimism on the fact that newer theories incorporate many theoretical constituents of their superseded predecessors, especially those constituents that have led to empirical successes. The substantive continuity in theory-change suggests that a rather stable network of theoretical principles and explanatory hypotheses has emerged, which has survived revolutionary changes, and has become part and parcel of our evolving scientific image of the world.

This reaction to the Pessimistic Induction has been initiated by Worrall's seminal (1989). What he called 'structural realism' is an attempt to capitalise on the fact that despite the radical changes at the theoretical level, successor theories have tended to retain the mathematical structure of their predecessors. Worrall's thought is that theories can successfully represent the structure of the world, although they tend to be wrong in their claims about the entities they posit. As we shall see in detail in the second part of the book, it turns out that this particular position is very difficult to defend (cf. Ladyman 1998; Psillos 1999, Chapter 7). Cartwright (1999, 4) has taken a different path. She is happy to go from the 'impressive empirical successes of our best physics theories' to 'the truth of these theories', but she denies that the assertions made by these theories are universal in scope. Rather, she goes for a 'local realism about a variety of different kinds of knowledge in a variety of different domains across a range of highly differentiated situations' (op. cit., 23) which tallies with her view that the world is best seen as disunified, with no laws or principles holding across the board and across different domains of inquiry. This is an issue that we shall discuss in detail in Chapter 6. Arguing as Cartwright does, for local truths which may vary from one model to another and from one domain to

another, may involve a perspectival notion of truth with characteristics not suitable for realism.

Realists talk of approximate truth and take science and its methods to issue in approximately true beliefs. How much of a substantive concession this is is a matter of dispute. Laudan (1984) claims that the realist cause is doomed unless a formal semantic for approximate truth is in the offing. Ron Giere (1988) concedes this but claims that realists can do well with a notion of similarity between the theoretical model and the domain to which it applies. Aronson, Harré and Way (1994) try to make good on the notion of similarity by devising an informal account of approximate truth which rests on the view that theories are type-hierarchies of natural kinds. Others (e.g., Niiniluoto 1999) still think that there are good prospects for a formal (and consistent) explication of approximate truth. My own view (cf. Psillos 1999, Chapter 11) has been that we shouldn't be deterred in our philosophical disputes by formal issues if the operative notions are intuitively clear and do not lead to paradoxes. As Peter Smith (1998) has suggested, the intuitive notion of 'approximate truth' can be explicated sufficiently well to be usable along the following lines: for a statement P, 'P' is approximately true iff approximately P. This relegates much to the concept of approximation, but there is no reason to think that a domain-specific understanding of approximation is not robust enough to warrant ascription of approximate truth in statements about each domain.

Although, as we have seen, there have been extremely important and profound challenges to realism, the only articulated rival philosophical position that has emerged is van Fraassen's (1980) *Constructive Empiricism*.[12] This view is already familiar to everyone and has been thoroughly debated in Paul Churchland and Clifford Hooker (1985). Its core point is that committed empiricists cannot be forced to be scientific realists because (a) they can offer an alternative account of science which takes science to aim at empirical adequacy and involves only belief in the empirical adequacy of theories; and (b) this account of science is complete in the sense that there are no features of science and its practice which cannot be accounted for (or explained away) from this empiricist perspective. Given that it is impossible to do justice to the massive literature on this subject in the present space (but cf. Rosen 1994 and Psillos 1999, Chapter 9), I shall only make a general comment on the spirit of van Fraassen's approach.[13]

As Richard Miller (1987, 369) nicely put it, van Fraassen's stance is a kind of modern 'principle of tolerance'. Although van Fraassen

(1980) can be easily interpreted as trying to show that scientific realism is an irrational attitude (and hence that constructive empiricism is the only rational attitude to science), in his later work (van Fraassen 1989, 1994, 2000a) he emphasises a new conception of rationality according to which constructive empiricism is *no less* rational than scientific realism. This new conception of rationality suggests that 'what is rational to believe includes anything that one is not rationally compelled to disbelieve' (1989, 171–2). Hence, van Fraassen says, since scientific realism is not rationally compelling, and since disbelief in constructive empiricism is not rationally compelling either, constructive empiricism is an equally rational option. All this may suggest that the door to scepticism is open, since from the fact that one is not rationally compelled to disbelieve P, it doesn't follow that one has (or could possibly have) good reasons to believe P. But van Fraassen (1989, 178) feels no threat here for he *denies* the 'sceptical' claim that 'it is irrational to maintain unjustified opinion'. This new aspect of van Fraassen's philosophy and his post-1990 attempt to articulate empiricism have not yet received the attention they deserve.[14] As an attempt to initiate this discussion, it might be possible to argue that there are tensions in van Fraassen's account of rationality. In particular, one could argue that from the fact that scientific realism is not rationally compelling it doesn't follow that constructive empiricism is no less rational an option. (Compare: from the fact that it's not rationally compelling to believe in Darwinism it does not follow that belief in Creationism is equally rational.) In order, however, to substantiate this tension, we need to show at least one of the following things. Either that there are aspects of the phenomenology of science which do not make good sense under Constructive Empiricism – for example, I think (Psillos 1999, 204) that the practice of diachronic conjunction of theories offers such a test-case. Or, that the joint belief in the existence of observable middle-sized material objects *and* unobservables is more rational than the combination of belief in middle-sized objects and agnosticism about unobservables. This last thought has been explored by Peter Forrest (1994). It's motivated by the claim that belief in the existence of unobservable entities (as opposed to agnosticism about them) rests on the same grounds as belief in the existence of middle-sized material objects (as opposed to agnosticism about them). This last claim, however, presupposes that there is no principled difference between having reasons to believe in the existence of observables and having reasons to believe in the existence of unobservables. Despite van

Fraassen's insistence on the contrary, there is a lot of sound philosophical argument that the equation of the unobservable with the epistemically inaccessible is bankrupt (cf. Churchland 1985; Salmon 1985).

Addendum

The preceding chapter aimed to offer a road map to the scientific realism debate. Since it was first published in 2000, a number of important books on scientific realism have appeared. Three of the most recent ones are Anjan Chakravartty's (2007), Derek Turner's (2007) and Christopher Norris's (2004) books. Taken together, these books present new challenges to realism and extend the debate to new territories. Here is some critical discussion of them.

A1. Semirealism: a short critique

A Metaphysics for Scientific Realism aims to do two things.[15] The first is to develop a viable realist position which capitalises on insights offered by entity realism and structural realism, while transgressing them. Semirealism, as Chakravartty calls it, comes out as a form of selective scepticism which restricts commitment only to those parts of theories that can be interpreted as describing aspects of the world with which scientists have managed to be in causal contact. The second aim is to develop a metaphysical framework within which his semirealism can be cast. This is a non-Humean framework based on a dispositional account of properties and a network of *de re* necessities. Chakravartty admits that this is just one option available to scientific realists, but claims that it gives semirealism a high degree of internal coherence, and hence facilitates its defence.

These two aims create a somewhat unstable mix. If semirealism is the best hope for scientific realists and if it is seen as *requiring* commitment to a non-Humean metaphysical picture of the world, this might be reason enough to make scientific realism unattractive to all those who prefer barren metaphysical landscapes. Semirealism is so much metaphysically loaded that its very posture might be enough to give extra force to well-known empiricist arguments that tend to favour antirealism on the grounds that it alone can deliver us from metaphysics. If this rich metaphysical picture is an add-on to the selective epistemic commitments of semirealism (if scientific realists do not have to buy it, anyway), why not leaving it behind, thereby making scientific realism a more inclusive philosophical position?

Indeed, Chakravartty focuses on the empiricist critique of metaphysics (advanced recently by van Fraassen) and contrasts van Fraassen's stance empiricism with what Chakravartty calls 'the metaphysical stance' (being taken to be largely the stance of scientific realism). Given van Fraassen's own permissive conception of rationality, the metaphysical stance cannot be shown to be incoherent; hence it cannot be shown to be irrational. So, Chakravartty claims, the empiricist critique of metaphysics cannot win: it cannot block realism from incorporating a rich metaphysical outlook. But then again on Chakravartty's set-up, realism cannot win either. At best, there will be a tie between the empiricist stance and the metaphysical stance.

The motivation for semirealism comes from the usual suspect: the pessimistic induction. This suggests that epistemic commitment should be restricted to those parts of theories that are more likely to resist future revisions. Semirealism adopts the epistemic optimism of entity realism (which is grounded on cases of experimental manipulation of unobservable entities), but adds that knowledge of causal interactions presupposes knowledge of causal properties of particulars and relations between them. Semirealism also adopts the epistemic optimism of structural realism (which is based on structural invariance in theory-change), but adds that the operative notion of structure should be concrete and not abstract. Concrete causal structures consist in relations of first-order causal properties, which account for causal interactions. Chakravartty claims that these causal properties are best seen as being powers, as having a dispositional identity. Focusing on *concrete* causal structures (with or without the power-based account of properties) is a step forward. It implies that one cannot have knowledge of the structures without also having knowledge of the intrinsic natures of things that make up the structure. Thus, Chakravartty claims, knowledge of concrete causal structures contains 'unavoidably' knowledge of intrinsic natures of particulars, and vice versa. This is all fine. It reveals some of the problems associated with structural realism. But, then again, why bother to call these things concrete causal *structures*? This is a term of art, of course. But in the context of the current realism debate, it is meant to imply a contrast between a concrete relational system and its (abstract) structure; it also implies that there can be knowledge of structural characteristics without concomitant knowledge of non-structural characteristics (e.g., of the entities that instantiate a structure). (For more on this, see Chapter 8.) Within semirealism, concrete causal structures (and their knowledge) contain everything up to the very natures of particulars. Since nothing is left

out, however, one cannot intelligibly talk about a substantial notion of *structure*.

Be that as it may, Chakravartty's key point is that the parts of theories to which realists should be epistemically committed should be those parts that can be interpreted as referring to a certain class of properties of concrete causal structures (or systems or whatever), namely, the 'detection' properties. These are properties that are causally detectable and in whose presence realists should most reasonably believe on the basis of the scientists' causal contact with the world. Detection properties are distinguished from auxiliary properties. These are attributed to particulars by theories but there is no reason to believe in their reality since they are not detected (though they might be detectable and become detected later on). Chakravartty appeals to what he calls a 'minimal interpretation' of the mathematical equations that make up a physical theory in order to demarcate the concrete causal structures associated with the detection properties from those associated with auxiliary ones. A minimal interpretation interprets realistically only those parts of equations that, in the context of a specific detection process, are indispensable for describing the (corresponding to that detection) concrete causal structures. Isn't there a tension here? If detection properties are specified independently of the theory, there is no need to interpret the theory *minimally* to get to them. If, however, they are specified in a theory-dependent way, this theory should already be interpreted prior to fixing the detection properties – and in all probability more than a minimal interpretation will be required to specify which properties are detection and which are auxiliaries.

As noted already, Chakravartty leaves the door open for less metaphysically loaded, but realist-friendly, conceptions of causation, laws and properties. His own view is that causation is a matter of continuous causal processes which are grounded in the dispositional nature of causal properties. Being powers, causal properties fix the laws in this and other worlds. They cast a net of *de re* necessities. This image of causal structuralism animates part of the book. Here is a worry, however. Semirealism has urged commitment to causally detectable properties and has clothed with suspicion all else (being merely auxiliary). But none of the extra stuff that Chakravartty (2007, 94) finds in the world (*de re* necessities, ungrounded dispositions and the like) are detectable. They are taken to be part of the baggage of semirealism because they play a certain explanatory role, notably they distinguish causal laws from merely accidental regularities (ibid., 94–5). But then, 'the deeper metaphysical foundations' of semirealism (ibid., 91) could well be (and

in all probability are) mere auxiliaries, which can then be treated with the same suspicion as other scientific auxiliaries (like the ether). This creates another tension. 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 based on the explanatory virtues of theories, in the case of the metaphysical foundations of semirealism, the only virtues on which one could base one's epistemic optimism are merely explanatory. Alternatively, if we allow explanatory considerations to play a role in science too (as distinct from mere causal detection), the very detection/auxiliary distinction that is so central in semirealism is put under a lot of strain. To put the point somewhat provocatively, the metaphysics of semirealism is the auxiliary system whose detection properties are Humean regularities and other metaphysically less fatty stuff.

A2. The natural historical attitude: a short critique

Turner (2007) takes the past to be epistemogically problematic in two important senses – both of which involve an epistemic asymmetry between the past and the present. First, there is the asymmetry of manipulability, namely, an inability to intervene (to manipulate) the past.[16] Second, there is the asymmetry of the role of background theories, namely, that background theories about the past imply (as background theories about the present and the tiny do not) that a lot of evidence about the past has been irrevocably destroyed and a lot of possible information channels have been dampened. These asymmetries make rampant local underdetermination of theories about prehistory. In a great deal of cases, we are told, scientists face an issue of choice between rival but empirically equivalent *and* equally epistemically virtuous theories, with nothing in their hands (no means to generate new phenomena, no sources of additional information) to break the tie. There is a third asymmetry, Turner notes, the asymmetry of analogy, namely, that past posits seem to be analogous to current observable entities. This, it might be thought, eases the problem of knowing the past. But Turner argues that it is precisely this asymmetry that explains why scientists have made a number of mistakes about the past. Hence, relying on analogy is not a reliable way to learn about the past. The general conclusion drawn from the three asymmetries is that there are clear senses in which knowing the past is harder than knowing the tiny and hence that scientific realism about historical sciences is in pretty

bad shape (and certainly in worse shape than scientific realism about electrons and genes).

Turner takes the past to be ontologically problematic too. There is a certain sense in which we should take the title of his book entirely literally: scientists do *make* prehistory. To be sure, that's not something Turner affirms. But he does not deny it either. He is neutral on this matter. It is a consistent hypothesis that the past is constructed and, for all we know, it might well be constructed. Turner's meta-philosophical stance is anti-metaphysical. Both realism and anti-realism (in all of its guises) impose metaphysical construals on scientific existential claims. They are not contended with saying that something exists or is real. Realists add a mind-independence gloss on existence/reality and anti-realists render existence/reality mind-dependent. Turner will have none of this. Does that remind you of something? Right guess! He is a fan of the *Natural Historical Attitude*, which is (a) an agnostic attitude towards metaphysical questions and (b) a 'gnostic' stance towards historical knowledge – we do have some knowledge of the past, though its extent should not be exaggerated, as the very idea of an epistemic access to the past is inexorably subject to the aforementioned asymmetries.

As noted already, Turner introduces an epistemic distinction between the past and the microphysical (the tiny). He claims that we can know more about the tiny than the past; hence, it is safer to be a scientific realist about the tiny unobservable. When it comes to the past, the defensible position is what he calls 'historical hypo-realism'. But are past things (e.g., dinosaurs) unobservable? Received wisdom has it that they are not – at least in the sense that they could be observed by suitably placed observers. Turner disputes received wisdom and claims that dinosaurs and their ilk *are* unobservable. Moreover, he argues that there are two distinct types of unobservable – the tiny and the past – and that this typical distinction bears an epistemic weight. Unobservables of type P(ast) are more difficult to be known than of type T(iny).

But is this quite right? Dinosaurs clearly are unlike electrons in terms of unobservability. The sense in which dinosaurs cannot be seen by naked (human) eye is different from the sense in which electrons cannot – different sort of modalities are involved here. Some laws of nature would have to be violated for either of them to be seen, but (interestingly) seeing dinosaurs (but not electrons) would not require a significant violation of the fundamental laws of nature. The possible world in which dinosaurs are observed is closer to the actual than the possible world in which electrons are observed. Observability is a matter of degree, but if we care to make a partition among the actually

observed, the observable and the unobservable, dinosaurs are closer to the middle than electrons. In any case, are there epistemically distinct types of unobservables? This cannot be an intrinsic difference of course; it will have to do with a principled difference between how an entity can be known by humans. Hence, it will have to do with the methods used to know that something is the case.

Carol Cleland (2002), based on the Lewisian thesis of the asymmetry of overdetermination, namely, that the present overdetermines the past but it underdetermines the future, has argued quite persuasively that there is, after all, an epistemic *symmetry* between knowing the past and knowing the present (and the future); hence, there is an epistemic symmetry between the methods of historical sciences and the method of physical (experimental, as she would put it) sciences. The idea is that historical scientists explore the present-to-past overdetermination to look for a tie-breaker between rival past hypotheses (a trace entailed by one but not by the other), whereas experimental scientists exploit the present-to-future underdetermination to devise experiments and establish predictions that can tell competing hypotheses apart. On Cleland's view, there is no principled difference between the two methods and both historical sciences and experimental sciences have an equal claim to justified belief and knowledge.

Turner disagrees with all this. What he offers as a reply, however, is not entirely convincing. He thinks (rightly) that there is widespread local underdetermination of past hypotheses. But he goes on to say that this kind of underdetermination is 'less common in experimental science' (Turner 2007, 57). Why? Because in historical sciences, unlike in experimental sciences, scientists cannot manufacture a crucial experiment. Asymmetry number 2 implies that we know that we have irrevocably lost crucial information about the past. There is, indeed, a difference here. But, is it not overstated? First, technological advancements (e.g., computer simulation) can provide plausible missing information about past processes. Second, the manipulation of the tiny can help break underdetermination ties only if we are allowed to bring into the picture the disparate theoretical virtues of competing theories. But we can do exactly that for competing hypotheses about the past.

Turner has interesting responses to this and other objections to his argument so far. For instance, he draws a distinction between a *unifier* (an entity that plays a unifying role) and a *producer* (an entity that can be manipulated to produce new phenomena), and argues that past (un)observables (like dinosaurs) can at best be unifiers whereas tiny unobervables can be producers too. On the basis of this he argues that

abductive arguments for past posits will be weaker than abductive arguments for tiny posits. Here is a worry, though. What about the past *and* the tiny, for example, a short-lived lepton? It is posited to explain (by unification, let us say) something that has already happened and though it is manipulable (in principle, at least) it was not manipulated in any way. Doesn't the same hold for a token of the type T-Rex? *Qua* a type of entity a lepton is both a unifier *and* a producer, though some tokens of it are posited as unifiers and others as producers (or both). The same goes for dinosaur (*qua* type): it is both unifier *and* producer, though (it seems that) all tokens of it are posited as unifiers and none as (actual) producers – though they did produce and they are subject to hypothetical manipulation.

Turner presents a number of distinct motivations for adopting forms of 'constructivism' – from Berkeley's, to Kant's, to Kuhn's, to Dummett's, to Latour's. Not all of them are, of course, constructivists and to the extent they can be lumped together as constructivists (anti-realists) their differences might be more significant than their similarities. The bottom-line of Turner's argument is that one may well remain agnostic on the issue of whether the past is real or constructed. This might be surprising for a reason that Turner does not seem to note. The discussion, in the bulk of the book, of local underdetermination, of the information-destroying processes etc. that are supposed to place the historical sciences in an epistemically disadvantageous position requires a sort of realism about the past. What sort? That *there are* historical matters of fact that would make one of the two (or many) competing theories true, but that somehow these facts cannot be accessed. If these facts are not independent in some relatively robust sense (an evidence-transcendent sense, at least), if they are 'socially constructed' as we go along, there is no reason to think that there will be (worse: there *must* be) significant gaps in the past. If facts about the past change over time, or if facts about the past are brought in and out of existence by scientists, it is not obvious to me why the past *resists* its incorporation into theories.

There is something quite puzzling in Turner's treatment of the issue of mind-independence. Here is (roughly) how he sets things up. Realists say: there are Xs (or X occurred), *and* they exist independently of the mental (or X occurred independently of the mental). Constructivists say: There are Xs (or X occurred), *and* they exist in a mind-dependent way (or X occurred mind-dependently). Given this set-up, he complains that the bits that occur after the 'and' are metaphysical add-ons to perfectly sensible scientific claims; they are not empirical hypotheses; they are not

confirmable by the evidence etc., etc. But these are not add-ons! Better put: they do not have to be construed as add-ons to perfectly legitimate empirical claims. Turner (2007, 148–9), to his credit, does note that this would be a natural reply, namely, that what is taken to be an add-on is really a way to unpack an existential claims (and there may well be different ways to unpack such claims). But what he says in reply (ibid., 149) seems to miss the point.

Turner's considered claim is that we are faced with a more general case of local underdetermination: scientific theories underdetermine the choice between realism (mind-independence) and social constructivism (mind-dependence). On top of this, we are told, there is reason to think that 'information about whether something happened mind-dependently or mind-independently will never get preserved in the historical record' (ibid., 156). I take it that Turner has some fun here. Too bad that universals, numbers, events and all the rest of the ontic categories do not leave any traces. If Turner is right, all metaphysics is killed off! Perhaps, Turner would be better off if he looked into the logical empiricist tradition of distinguishing between empirical realism and metaphysical realism. One may well be able to leave metaphysics behind without simply being neutral on the realism–constructivism issue.

The Natural Historical Attitude (NHA) that Turner defends is partly an antidote to constructive empiricism. Turner argues (following Kitcher) that constructive empiricism implies scepticism about the past. This, however, depends on whether we think that past posits are observable or not. And we have discussed that already. It is welcoming news that NHA (like its parent, the Natural Ontological Attitude) is not a sceptical stance. The key idea behind this anti-sceptical stance is what Turner calls 'the Principle of Parity' (what I have called the no-double-standards principle), namely, the very same methods of confirmation apply to claims that purport to refer to both observable and unobservable entities and hence that claims about unobservables can be as well supported by the relevant evidence as claims about observables.

A3. Constructivist anti-realism: a short critique

Christopher Norris's (2004) book fills an important gap in the debate on scientific realism by looking into Norwood Russell Hanson's philosophy of science.[17] Though Norris does not quite put it that way, I think Hanson's key (though neglected before Norris's book) contribution to the debate was that he made possible a non-sceptical version of scientific anti-realism. Put in a nutshell, Hanson's idea is that one can believe

in whatever is posited by modern science *and* accept that theories are, by and large, true, while avoiding metaphysical commitments to a mind-independent world and robust realist accounts of truth. The details of Hanson's position are intriguing and subtle and Norris does an excellent job in describing (and criticising) them. But it is useful to keep the big picture in mind, if we are to assess Hanson's contribution. Hanson died prematurely and we can only speculate as to how his views might have developed. But, as Norris amply illustrates, he felt the tensions and problems of the position he was trying to develop.

Hanson was deeply influenced by the later Ludwig Wittgenstein, and in his turn, he deeply influenced Thomas Kuhn and Paul Feyerabend. He employed centrally the Wittgensteinian idea that there is not a ready-made world. Rather, what there is and what one is committed to depends on the 'logical grammar' of the language one uses to speak of the world. Wittgenstein's 'logical grammar' was meant to capture the interconnections of the uses of key concepts that structure a certain language-game. Science is no less a 'language game' than others. This game is characterised by its norms, rules, practices and concepts, but all these are *internal* to the game: they do not 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 (whatever that is) and come to accept theories as true as well as believe in the existence of unobservable entities. One, that is, can behave as a scientific realist: one need not be a sceptic. But, on Hanson's view, one need not (perhaps, should not) add to this behaviour any robust realist metaphysics. Nor should one build into the language-game a concept of truth that is evidence-transcendent.

Hanson's philosophy of science has three important entry points. The *first* comes from the rejection of the empiricist view that there can be a theory-free observational language. In fact, Hanson went much beyond this negative thesis. Based on Wittgenstein's claim that all seeing is 'seeing as', he argued that all perception is aspect-relative: there is no way in which Tycho Brahe and Kepler saw the 'same thing' when they turned their eyes to the heavens, since eyes are *blind* and what they see depends on what conceptual input shapes their seeing. This positive thesis leads quickly to claims of incommensurability. In fact, as Norris points out, it leads quickly to perceptual relativism, which renders impossible any attempt to make sense of the empirical basis of science in a way independent of the language-game we adhere to. This first entry point loses the world as a mind-independent structured whole, but re-instates a

paradigm-relative world, namely, a world of the phenomena as they are shaped by a certain language-game.

Hanson's *second* entry point comes from quantum mechanics. He bought into the orthodox (Copenhagen) interpretation of it and thought that this leads to inevitable changes in the way we see the world and the way we raise epistemological questions. Presumably, quantum mechanics reveals the inherent limitations of the claim that the world is objective and mind-independent. It also sets limits to what can be known of it and to what kind of theories can be true of it. Norris discusses in some detail Hanson's disapproval of any Bohm-like theory of quantum phenomena and his commitment to a radical discontinuity between the quantum world and the classical one. Here again, Hanson drew the conclusion that accepting quantum mechanics (in its orthodox interpretation) amounts to adhering to a new language-game in light of which the old (classical) language-game makes no sense.

Finally, Hanson's *third* entry point comes from his work on the 'logic of discovery'. In his *Patterns of Discovery* (1958), he did perhaps more than anyone else to legitimise abduction, namely, the mode of reasoning according to which a hypothesis is accepted on the basis that it offers the best explanation of the evidence. Hanson was no friend of instrumentalism. He had no problem with taking scientific theories at face-value. He had no qualms about scientists' going beyond the observable evidence and accepting the existence of unobservable entities on an abductive basis. These unobservable entities are neither logical fictions nor merely hypothetical. They are part of the furniture of the world. But of which world? Hanson's answer is again tied to the idea of language-games: the world as specified by the language-game of science. This world is infested with causal-explanatory connections that underlie the legitimate uses of abductive reasoning, but these connections are, again, the product of several linguistic rules and practices.

It is not hard to see how these three entry points make possible a non-sceptical version of scientific anti-realism: science is not in the business of discovering the structure of a mind-independent world. Rather, it is the language-game that imposes structure onto the world and specifies what facts there are. Accordingly, science can deliver truth, but the truth it does deliver is determined by the epistemic resources, practices and norms of the language-game that constitutes science. This is more evident, as Norris (2004, 113 & 115) notes, in Hanson's notion of a *pattern*. Patterns (the ways objects are conceived) have empirical implications but they are *not* themselves empirical: they are imposed by the conceptual scheme and to deny them is to attack the conceptual

scheme itself. There is a Kantian ring in this view. But in Hanson's case, the result is relativised Kantianism. For, in light of perceptual relativism and incommensurability, there is a plurality of language-games each of which constitutes its own phenomenal world. The 'objective' world is either lost or reduced to a *noumenal* blob.

This last point brings to the fore a central problem that Hanson's anti-realism faces: how is change explained? To his credit, Norris (see especially 2004, 37–9) makes capital on this problem on behalf of scientific realists. Here is how I would put the matter. Hanson's view comes to this:

Constitution: The worldly objects that science studies are (conceptually) constituted as objects by the language-game (conceptual scheme, rules, theories and practices) that scientific theories use to study the world.

This thesis, however, is in tension with an empirical fact:

Refutation: The conceptual schemes that science uses to study the world are revis*able* and revis*ed*.

If scientific objects were constructed/constituted by the conceptual resources of theories, would it not be natural to expect that the very same conceptual resources of the theory would be able to constitute all relevant objects? In particular, how can there be friction *within* the conceptual scheme? Would any friction be either impossible or else explained away by the right constitution of objects? Why, for example, if the relevant scientific objects are constituted by Tycho Brahe's framework, should some phenomena lead scientists to *abandon* this framework? There is a very strong intuition, I think, that the friction can only be explained if the world (something external to the conceptual scheme) exerted some *resistance* to our attempts to conceptualise it.

This intuition, together with *Refutation*, might be thought enough to refute *Constitution*. But there seems to be a way out for its advocates. It might be argued that the world is indeed there, but only as a structure-less (or minimally structured) mould. Yet, this is no improvement. Suppose that the world is a structure-less (or minimally structured) mould. We know that the presence of anomalies to scientific theories is *diachronic*. Anomalies do not go away too easily. Sometimes several modifications of our current theory/conceptual scheme have to be tried before we hit upon the one that removes the anomaly. Besides, anomalies do occur in the theories/conceptual schemes that replace the existing ones. If the world were merely a structure-less mould, then this recurring friction could not be explained. A structure-less mould can

be shaped in any way we like. And if it is shaped in a certain way, there is no reason to expect that the shaping will turn out to be inadequate, unless the mould has already, so to speak, a *shape* – a natural causal structure. If the world has a certain causal structure, it is easier to explain why some attempts to fix an anomaly are better than others, as well as to explain why some such attempts prove futile. Hence, if we allow the world to enter the picture as the explanation of friction (and of the subsequent replacement of one's preferred phenomenal world), we'd better also think that this world has already built into it a natural causal structure.

Index