Quests of a Realist

Stathis Psillos, *Scientific Realism: How Science Tracks Truth*. London: Routledge, 1999. Pp. xxv + 341. £16.99 PB.

By Michael Redhead

This book provides a carefully argued defence of scientific realism. It is attractively written, and demonstrates that analytic philosophy does not need to be presented in a turgid, impenetrable style. It is very much a product of the 'London' style of doing philosophy of science with a nice balance of logical precision and sensitive historical case studies, and is also very up-to-date in its coverage of recent literature on this much debated and highly contentious topic.

Psillos makes great play with one of the main arguments for realism, the 'No Miracles' argument. According to this argument the crucial feature of the mature sciences is not just to codify existing knowledge, but to make novel predictions that, if verified, produce abundant confirmation for the reality of the underlying entities and processes responsible for these predictions. If the theoretical discourse was just a fiction it would be a miracle if science was able to make empirically successful novel predictions, and it is simply not rational to entertain any account of science which relies on miracles to explain its success. All this sounds plausible, but needs careful discussion of what we mean by novel prediction and just how this relates to confirmation.

Psillos distinguishes what he calls 'use novelty' from temporal novelty. A known observational fact *e* is use-novel with respect to a theory T if no information about this fact was used in the construction of the theory, which nevertheless 'predicts' *e*. By contrast a fact *e* is temporally novel with respect to T if it was not known before T was proposed. Clearly temporal novelty implies use novelty, but the converse does not hold. The classic example of use novelty is the motion of the perehelion of Mercury in relation to the general theory of relativity (GR). This astronomical fact was of course known before Einstein proposed GR, but its 'prediction' by GR seems to offer genuine support for the latter theory. Psillos supports this view but claims that temporal novelty should carry additional weight (p. 107), but then in a seemingly contradictory footnote says that, after all, comparative confirmational weights can only be settled by close examination of actual scientific practice.

The trouble is that Psillos gives no adequate explication of what he means by confirmation. As an example that goes against Psillos let us consider one version of a Bayesian account (for other versions see Howson and Urbach [1993]).

For Bayesianism it is often claimed that known facts can never provide confirmation for a theory T that predicts them. This is based on the Bayesian rule

$$\text{Prob}\ (\text{T}/e\ \&\ b) = \text{Prob}\ (\text{T}/b)/\text{Prob}\ (e/b)$$

where b is background knowledge. Trivially, if *e* is part of b then Prob (*e*/*b*) = 1 and so the posterior probability Prob (T/*e* & *b*) is equal to the prior probability Prob (T/*b*) and *e* has provided no incremental confirmation for T. This goes under the name of the 'problem of old evidence'.

But this problem can be handled by the Bayesian in the following way. Baysianism allows no discrimination between hypotheses with the same empirical adequacy, other than what is loaded into the priors. Essentially what temporally novel evidence achieves in Bayesian conditionalisation is to reduce to zero the probability of any hypothesis that predicts not-*e* and to redistribute the probability measure amongst all hypotheses that predict *e* in exactly the same proportion as obtains with the priors.

Consider now the case of ad hoc accommodation, where a theory is 'designed' to explain *e*. Surprisingly enough, Bayesianism allows that *e* can still provide support for T, indeed just as much as in the case of temporal novelty.

Thus if *e* is used heuristically as a 'filter' on theory construction, then any hypothesis that predicts not-*e* will be filtered out and accorded zero prior probability, the whole probability measure being spread amongst those hypotheses that do predict *e*, in proportion to their priors as assessed before filtering. But this just means that the priors after filtering are exactly equal to the posteriors after Bayesian conditionalisation in the case of temporal novelty. Of course, after filtering, the posteriors will equal the priors, so there is no additional confirmation as we have already remarked, but the posteriors in the two cases of temporal novelty and ad hoc accommodation are numerically identical! Use novelty does not enter into the Bayesian account, if the scientist is assumed to be logically omniscient so that coherence forces her to use *e* as a factor in assigning the prior probability of any theoretical proposal. Of course the assumption of logical omniscience can be criticised as unrealistic and other accounts of confirmation deliver different verdicts. But since Psillos attaches such importance to novel predictions in support of realism, it is a pity that he makes no attempt to underpin his view with some formal account of confirmation.

Another problem that arises here is the problem of underdetermination: there may be many theories, perhaps even an infinite number, that have the same empirical content, and hence empirical considerations can offer no help in deciding between them. How then should the scientist proceed? This is where the pragmatic virtues of simplicity, unifying power and so on, come in. Crudely, the scientist chooses the simplest hypothesis that will explain the evidence. But simplicity is a tricky notion to explicate formally. It certainly has something to do with the paucity of adjustable parameters incorporated in the hypothesis, and hence is connected with Popperian themes of non-adhocness and testability, but in the advanced mathematised sciences, it is related to more intuitive ideas about the nature of the specific mathematical apparatus such as differential equations used to represent physical phenomena.

But even if we have decided what constitutes a simple unified theory, is this a guide to truth? In Bayesian terms, should simpler hypotheses be accorded higher priors? There are two ways to go here. Firstly, we may define a concept of scientific rationality as one which invokes the simplest, most unified theory, to explain empirical phenomena. This argument in defence of the scientific account is by itself clearly viciously circular. Psillos himself is sympathetic to a second approach, namely that the past record of scientific theories, in producing successful novel predictions, for example, can be used to justify the pragmatic explanatory virtues. In the past science has used the pragmatic criteria, and as a matter of historical fact, this has led to progress at the level of empirical adequacy. So is it not natural to expect the same criteria to produce more successful science in the future? Such meta-inductions are always liable to fallibility. Perhaps at some deep level of explanation physics, for example, will just get more complicated rather than increasingly simple. The pragmatic virtues are best thought of perhaps as a regulative ideal which guides the scientific enterprise, has been successful at generating progress at the empirical level, particularly as marked by successful novel predictions, but which may not be indefinitely achievable. This approach via the notion of a regulative ideal combines aspects of both approaches to the pragmatic virtues we outlined above.

But talking of meta-inductions leads to another famous difficulty for realism, the argument of the so-called pessimistic meta-induction. This recognises that the progress of science is not a monotonic process, but is punctuated by revolutions in which most, if not all, the old ontology is simply abandoned. So if realists were justified in accepting entities like phlogiston or the luminiferous ether, and these concepts were later jettisoned, should we not conclude that all the present-day theoretical talk of quarks and photons and so on is no more than fiction, the only serious basis for appraising science is at the observational, not the theoretical level?

One response to the pessimistic meta-induction is to try and show more theoretical continuity across scientific revolutions than the potted versions of the history of science indulged in by philosophers of science typically allow. This is essentially the tack taken by Psillos. He rejects the simplistic causal theory of reference for theoretical terms, which would too trivially guarantee referential stability across theory change, for example allowing that proponents of the phlogiston theory of combustion were really referring to oxygen all the time, given that oxygen rather than phlogiston, as we now believe, plays the true causal role in combustion phenomena. Psillos follows instead a causal-descriptive theory of reference in which a core set of properties is associated with the causal agent, in addition to the mere fact that it is the causal agent, and referential stability is now linked to preservation of those core properties. In this account Psillos argues for referential stability of caloric and ether, but not of phlogiston. But the discussion looks not so much like philosophical analysis, but rather involves peering into the psychology and/or private notebooks to ascertain what scientists really meant by terms like 'ether' or 'phlogiston'. Psillos presents detailed case studies for the examples of caloric and ether but what the discussion boils down to seems to be that structural aspects of the old theory are preserved in the new theory.

This suggests a move to Worrall's (1989) structural realism programme, in which what is preserved across theory change are mathematical structures as captured by the equations governing the behaviour of the posited theoretical entities, rather than the intrinsic nature of the entities themselves. Psillos does, in fact, devote a whole chapter to Worrall's views, but while admitting the importance of some of Worrall's ideas, he ends up dismissing the programme in favour of the causal-descriptive analysis we have outlined above. In my view Psillos does not properly do justice to Worrall's ideas, but furthermore Worrall can be criticised, though for reasons different from those put forward by Psillos. So let me now make some comments on the pros and cons of structural realism.

Firstly, Psillos never explains clearly what he means by structure. The vague reference to mathematical equations is not sufficient to get a proper handle on this notion. Informally a structure is a system of related elements, and structuralism is a point of view which focuses attention on the relations between the elements as distinct from the elements themselves. So think of building materials that fit together to make a house, or brush strokes which relate to form a picture, or words which string together into meaningful sentences.

All these are examples of concrete structures. But to define an abstract structure we can imagine collecting concrete structures into isomorphism classes, where two concrete structures in the same isomorphism class are related by a bijective correspondence which preserves the system of relations in the sense that if in the one structure the elements $x_1, x_2 \ldots x_n$ satisfy the n-ary relation R, then in the second structure the corresponding elements $y_1, y_2 \ldots y_n$ satisfy $R' (y_1 \ldots y_n)$ if and only if $R (x_1 \ldots x_n)$, where R' is the n-ary relation in the second structure which corresponds to R in the first structure.

The concept of abstract structure can now be thought of in an *ante rem* Platonic sense as the second-order Form which is shared by all the concrete relational structures in a given isomorphism class; or in extension we can conceive of the abstract structure just as the isomorphism class itself, which can be represented by any arbitrarily selected member of that class. In particular, mathematical structures invoking, for example, natural numbers or real numbers can be used to represent the abstract structure associated with a physical system if they belong to the same isomorphism class. The claim of the structural realist is that this abstract structure associated with physical reality is what science aims, and to some extent succeeds, to uncover, rather than the true physical relations of that reality. The abstract structure can be thought of then as a second-order property of the true physical relations, rather than these physical relations themselves.

The mathematical representation of this abstract structure is what Psillos and Worrall mean when they talk about the 'equations of a theory'.

Psillos seems to assume that the structural realist is committed ontologically only to the reality of abstract structure. He then proceeds to demolish this view, in particular by citing the famous result of Newman (1928) that the Ramsey sentence $\exists R(S(R))$, asserting the existence of a relation R which has structure S, is in fact a logical truth, modulo the specification of the cardinality of the domain over which the relation is defined. So the structural characterisation of the relation is essentially contentless!

But this pure ontological version of structuralism is not necessary to understand the Worrall program. It is a matter of epistemology rather than ontology. We need not deny that there are real physical relations posited by physical theories. The question is how much can we reliably claim to know about these relations. The pessimistic meta-induction suggests we never know the relations themselves, but Worrall's claims about the stability of structure prompts the view that it is just the structural features (as represented for us by the mathematical equations, crudely speaking) that can be judged as true knowledge about the external physical reality. Thus S(R), where R refers to a specific relation having the structure S, is of course logically stronger than the Ramsey sentence, and is by no means a logical truth. But this means, as Psillos rightly emphasises, that the reference of R must be picked out in non-structural terms. But this is not denied in the above account. Our claim is merely that R is hypothesised in some explanatory theoretical context so it exists as an ontological posit, but all that we have epistemic warrant for is the second-order structure S.

Psillos continues by arguing that the distinction between nature and structure is unclear. Surely part of what we mean by the nature of an entity is the structural property of the relations into which it enters. I don't at all disagree with this point. But this is really a semantic red herring. All that the structural realist needs to claim, on my account, is that part, i.e. the structural part, of the nature of the posited physical entities is all that we can claim to know.

So far I have been defending Worrall against Psillos, but now I would like to make some criticisms of my own.

Although a number of historical episodes of theory change seem to fit the Worrall model, for example the transition from Fresnel to Maxwell in the theory of the reflection and refraction of light at the interface between two dispersive media, which is discussed in some detail by Psillos (with the opposite conclusion!), nevertheless there are many counter-examples at any rate to a simple version of the thesis. Consider the case of classical neo-Newtonian spacetime being replaced by the Minkowski spacetime of special relativity. We can consider a family of structures $\{S_c\}$ corresponding to varying the velocity of light c. For all finite c we can argue that the structure is stable with respect to changing c, but for $c = \infty$ there is a qualitative singularity in the sense that the metric of spacetime becomes singular in this limit. The existence of qualitative singularities of this type is also apparent in the case of the family of quantum-mechanical structures indexed by a variable Planck's constant h. $\{S_h\}$ is structurally stable for all values of h unequal to zero, but for $h = 0$ the family of structures again exhibits a qualitative singularity, in the sense that the noncommutation property of appropriate quantum-mechanical observables (strictly of their associated self-adjoint operators) is abruptly and discontinuously lost. But, of course, the singular transition from $h = 0$ to a finite h, or from an infinite c to a finite c are exactly what characterises the conceptual revolutions engendered by quantum mechanics and relativity. It is the singularity in the family of structures $\{S_h\}$ and $\{S_c\}$ which marks the really revolutionary aspects of the new theories, and here it just seems wrong to claim that the mathematics has survived qualitatively intact. These are the cases where Worrall's account seems to break down.

Of course one can claim that the move from a singular metric to a non-singular metric in geometry or from a commutative algebra to a non-commutative algebra in the quantum-mechanical case are heuristically natural moves to make qua pure mathematics, so that one may want to argue that physics grows by a process of 'natural' modification of mathematical structure, but that is quite different from the claim that there is continuity 'at the level of the equations'. So the structural realism programme seems to me to face problems, at any rate in any simplistic version that makes no appeal to the 'natural' growth of mathematical ideas in physics.

Centre for Philosophy of Natural and Social Science,
London School of Economics and Political Science,
London WC2A 2AE, UK.

By Peter Lipton

Stathis Psillos has given us an exceptionally rich and vigorous defence of the idea that science is in the truth business. At its heart lies what Psillos refers to as the 'No Miracles' argument. The motivating idea is that the predictive successes of our best scientific theories gives us a reason to believe them (approximately) true, and a reason to believe that the methods that led scientists to accept these theories are reliable tools for discovering the truth, thus a reason for scientific realism. In slogan form, the reason for realism is that it is the best explanation of scientific success. Psillos defends this argument against two central objections. The first is that it is viciously circular, since in effect it uses inference to the best explanation to justify inference to the best explanation. The second objection is that in any case the No Miracles argument fails on its own terms, since the truth of theories is not the best explanation of their predictive success. It is these two objections and Psillos's responses to them that will be the focus of my comments.

Let me begin by saying a bit more about the form of Psillos's version of the argument. However poorly philosophers may understand their practices, scientists have ways of determining when a scientific theory should be accepted as true. Psillos and I both think that abduction or inference to the best explanation is an important part of the story, but in what follows I will typically use these expressions as a way of referring to scientists' inductive practices in general, whatever form they actually take. Using predictive success as a guide to inference will clearly be an important element of these practices, but it is not the whole story, since one theory may be judged to be more belief-worthy than another, even when both are successful. Let us call these other factors that make a theory belief-worthy the epistemic virtues. Psillos singles out two virtues in particular. One is that a theory plays an indispensable role in successful predictions, where this means that without the theory the inference to the prediction would no longer be deductively valid, and that there is no other sensible theory that could replace the first to restore validity (p. 110). The other central virtue is that some of these successful predictions should be novel (pp. 105–107).

The central realist claim is what I will call the 'realist theory', that scientists' inductive practices are reliable routes to the truth: these practices tend to take scientists to theories that are true. The No Miracles argument then looks like this. We have reason to believe the realist theory because, by the standards of scientists' own inductive practices, it is well-supported by the continuing predictive success of the scientific theories to which those same practices have led. These theories have been retained as science develops, and this retention—due primarily to continued predictive success—is the best reason we could have to believe the realist theory, because the realist theory is the best explanation of that retention.

I turn now to the vexed question of the circularity of the No Miracles argument. According to the circularity objection, even if the truth of a scientific theory would be the best explanation of its predictive success, that provides no reason to believe the theory is true, since the reliability of inference to the best explanation, the method by which the theory was inferred in the first place, is precisely what is at issue. Psillos's reply is that there is no vicious circle here, since a rule of inference can be used without being assumed and will provide justification just in case it is in fact reliable.

My own reactions to the circularity objection and to Psillos's reply are mixed. First, the objection does alas appear to show that the No Miracles argument preaches only to the converted: it has no probative force for those who are not already inclined to use inference to the best explanation. Psillos seems to agree (cf. pp. 88–89), but finds this much less depressing than I do. What explains this difference? Psillos's view is that our inductive practices are basic for us, so no justification that does not use them is either possible or required. But not required for what? Clearly not required in order for it to be possible that induction be in fact reliable, and perhaps not even required to provide some kind of justification for that claim. But an independent justification is required to answer the radical sceptic who will not begin by using induction. To call a practice basic is another way of saying that such a sceptic cannot be answered.

So even if one grants that induction is basic for us, this does not in my view defuse the radical sceptical hypothesis that induction is unreliable even though it passes all its own tests. Moreover, it does not appear that our inductive practices, as the realist construes them, are in fact basic for all of us. For there are, alas, enemies of realism. These benighted souls may have to indulge in some kind of nondemonstrative reasoning, but it may not be the kind that supports the No Miracles argument. Indeed some anti-realists may endorse reasoning of the same form as the realist—inference to the best explanation, if you like—but they do not endorse the claim that the conclusions of such inferences are true, in the realist's sense. They may be constructive empiricists, who only infer empirical adequacy, or they may be Kuhnian Kantians, who claim that the world described by a scientific theory is one partially and variably constituted by the conceptual equipment of the scientists who study it. Because the No Miracles argument assumes the realist form of induction that it seeks to justify, it appears that the argument will be ineffective against these opponents.

My second reaction is more positive. It is to agree with Psillos that those of us fortunate enough to use induction may usefully deploy inductive arguments to assess the reliability of inductive methods or practices, however impotent we may be against the sceptics and the anti-realists. As I see it, this possibility of internal justification is tantamount to the obvious possibility of testing the reliability of instruments. An instrument is from this point of view like an inductive rule, allowing us to go from the premise that, say, the barometer is falling, to a conclusion that is not deductively entailed, say that there will soon be a storm. When we assess the reliability of this rule, by seeing how frequently it works, we are providing an inductive assessment of an inductive rule. Psillos gives the example of an 'inference machine' that makes predictions from the premises we provide as input (p. 84). The observed track record of such a machine would provide inductive evidence of the reliability of the inductive rules it follows. Moreover, an inductive assessment of an inductive method may be legitimate even when the general method being assessed is the same as the method used to do the assessing.

So my second reaction is that there can be legitimate inductive assessments of inductive methods. But my third and final reaction is now to question whether the No Miracles argument is one of these. In brief, my worry is that, unlike the clearly benign cases of the inductive justification of an inductive practice, the No Miracles argument appears to introduce no new evidence for the truth of successful theories and so no new evidence for the reliability of inference to the best explanation as a route to true theories. To see the problem, consider first a legitimate case. Insofar as barometers are reliable predictors of storms, when they fall they give me reason to believe there will be a storm. From the externalist perspective that Psillos and I share, these reasons do not depend on checking the reliability of barometers: it is just the fact of their reliability that matters. But before I do check their track record, I have no reason to say they are reliable. After checking, however, I have new evidence, namely the correlation between the predictions and the actual weather, and this gives me a justification for a claim about the barometers' reliability. In the case of the No Miracles argument, by contrast, there seems to be no new evidence. The predictive successes of a scientific theory provides the scientist with evidence for its truth; but when the philosopher comes on the scene *she* does not gather further evidence of the track record of particular theories. Rather, *she* simply says that when theories are successful they are probably approximately true. But the successes themselves are already part of the scientific case for the theories that enjoy them.

Is there any way to construe the No Miracles argument so that it does introduce, or at least draw attention to new evidence? Psillos's discussion suggests two possibilities. Psillos is impressed by the way an inferred theory may be retained as new evidence comes in and as new tests are conducted. Our inferential practices sanction certain theories; then these theories go on to have further successes. It is the continuing success that provides new evidence and thus a special reason to believe that those practices are reliable. The second way of trying to construe the No Miracles argument as introducing new evidence is to focus on the credit that predictive success confers on background theories. Following Richard Boyd, Psillos's version of the No Miracles argument suggests that the success of a scientific theory under test reflects well not just on that theory, but on all the other theories that are assumed as part of the essential background to the testing process (cf. p. 78). So one may then be able to argue that the No Miracles argument is bringing more evidence to bear on a given theory than the first-order scientific case, since we now have both the evidence from the tests a theory has passed and the evidence that supported other theories where the first theory served as an essential part of the background.

These moves give us more sophisticated versions of the No Miracles argument, but they do not seem genuinely to add new evidence for the truth of scientific theories. The reason is that the continued success of our theories, both in foreground and background, are in the end also part of the first-order, scientific case for their truth. Of course a scientist might not notice say that a theory was tacitly assumed in a test of another theory, and so not notice that the assumed theory also received credit from the success of the test. But this would still be part of the scientific case, even if a philosopher were to draw attention to it. And the scientific case for our best theories is what the No Miracles argument assumes, not what it establishes.

Perhaps Psillos is not claiming that the No Miracles argument provides any evidence for the truth of scientific theories beyond the scientists' evidence. He does insist that the argument goes beyond scientists' own inferences on the grounds that while scientists claim that particular theories are true, the No Miracles argument defends the realist theory—the general claim that science's methods are reliable routes to the truth (p. 79). But this insistence may not be on new evidence. It may be rather be on the point that the conclusion is different in the two cases, the philosophers' inferred from the scientists'. Here it seems that scientists tell us that particular theories are true, and philosophers use this result to go on to show that the scientists' methods are reliable.

On this construal of the No Miracles argument, the realist's argument is indeed different from the scientists', but we have a tighter circle than we may have at first realised. It is not simply that the realist is using the very form of inference she seeks to justify: she is also using the output of scientists' inferences—the claim that their best theories are true— as a premise of her argument, from which she infers the general reliability of their practices. This argument may be unexceptionable, but it makes it clear that the realist here is in no way introducing new evidence or testing the scientist's methods. Rather she is simply moving from the assumption that those methods have worked in the past to the conclusion that they will work well in general. What we have in the end, it seems, is the drawing of a general moral from the prior commitment to the truth of specific theories. That moral may be worth drawing, but it is not much of an argument for realism or, to put is slightly less negatively, it is not much more of an argument for realism than the scientific case for the truth of particular scientific theories.

The other central challenge to the No Miracles argument is that the realist theory that our practices are reliable routes to theoretical truth is not in fact the best available explanation for the predictive successes our best theories have enjoyed. Psillos considers several alternative explanations for predictive success and retention and his critiques of these alternatives are in my view telling, but there are two other competitors that I will briefly consider.

The first of these is really a class of competitors, the 'underdetermined alternative theories'. Given any successful theory, there are always other, competing theories that would have enjoyed the same successes: these are the underdetermined alternatives. The question then is why the best explanation for the success of a particular theory is the truth of that theory, rather than the truth of any of these alternatives. The second alternative explanation appeals to Kuhn's conception of scientific achievement. As I read him, Kuhn holds that scientific theories do describe the world, but this world is partially constituted by conceptual activities of the scientists. This picture allows for a notion of truth within a single normal science tradition, roughly that of correspondence to the phenomenal world in which that tradition operates.

The underdetermined alternatives and the Kuhnian alternative provide foils to the realist theory at different levels. The former accepts a realist understanding of truth, but asks why we should suppose that it is the theories that scientists' practices take them to that are the true ones, rather than others that would have enjoyed the same sort of success. The Kuhnian alternative does not propose different theories, but rather a different conception of truth. But both foils offer alternative explanations for predictive success. Why then is the realist theory the best explanation?

Psillos's response to the underdetermined alternatives would I think be to deny that these alternatives are equally good explanations of the success they share with the theory the scientists actually accept. They are not equally good because they do not display various epistemic virtues to the same degree. This appeal to virtues underlies Psillos's response to the underdetermination of theory by data. The virtues provide the extra constraints that may eliminate the underdetermination.

What about the Kuhnian alternative? An appeal to the virtues does not seem helpful here, since the Kuhnian would have us infer the same theory as the realist, only the Kuhnian gives that inference a different gloss. What the defender of the realist theory can say, I suppose, is that the list of virtues, and indeed the account of inductive practices generally, only seem to be shared, because while they may have the same form, their content is different, simply because scientists' actual practices are, according to the realist, to make inferences to realist truth, not to Kuhnian truth.

Supposing that the realist succeeds in this strategy of product differentiation, she may then be in a position to claim that, given her list of virtues, the realist theory is indeed a better explanation of predictive success than either the underdetermined alternatives or the Kuhnian view of truth. But this leaves me with two final worries. First, what is the status of the justification the realist has for her list of virtues? For Psillos that justification is part and parcel of the No Miracles argument. But if as I fear the realist is simply assuming that the scientists' practices have taken them to the truth, then it seems that the foils to realism are being excluded without reason; if she does not so assume, then it is unclear that the foil explanations are worse than the realist theory.

My other worry is this. Even if we take it that the epistemic virtues are justified by the predictive success of the theories that have them, there remains the different question of whether the realist theory itself enjoys these virtues. The realist theory is that our inductive practices are a reliable route to the truth. Is that indispensable to the prediction of the evidence? I cannot see that it is. For the evidence is that that our best theories are successful, and while that is entailed by the realist truth of those theories, it is also entailed by their empirical adequacy, by the truth of any of their underdetermined alternatives, and I suppose by their Kuhnian truth. What about novel predictions? It is not obvious what these might be. Particular scientific theories sometimes make novel predictions, but what novel prediction follows from the claim that scientific practices can attain the truth?

The point here is not that there are better explanations of success about; the alternatives I have mentioned do not possess the virtues of indispensability or novel predictive power either, so far as I can tell. But our own inductive practices require before we infer an explanation not merely that there be no better one available, but that the one at hand be good enough. So if the realist hypothesis lacks two central epistemic virtues, by the realist's own standards it ought not to be inferred. Perhaps the correct moral of this story is rather that the realist theory is not really on a par with scientific theories and so should not be judged in the same way. But that route, I suspect, will just take us back to the position that, at the end of the day, the only evidence for the truth of scientific theories is the evidence that scientists' use, and the only positive arguments for scientific realism are the arguments that scientists make.

Department of History and Philosophy of Science,
Free School Lane,
University of Cambridge,
Cambridge, UK.

By Igor Douven

Psillos' strategy is two-pronged: on the one hand, he tries to defend scientific realism 'head-on' by greatly elaborating on the well-known explanationist defence of scientific realism initiated by, among others, Putnam and Boyd. On the other hand, he tries to defend scientific realism by arguing that, in the end, we are unable to make sense of the alternatives to it that have been proposed, whether these are of a realist variety like, for instance, structural realism or of an antirealist variety, such as van Fraassen's constructive empiricism and Fine's NOA. Here I will mainly concentrate on the first part of the strategy, though this will also lead me to say something about Psillos' critique of van Fraassen's position. I should say right away that I have great sympathy for the way in which Psillos seeks to defend realism and that, in my opinion, his defence goes further and is better than anything realists have said so far in support of their position. Still, Psillos' defence may give rise to some worries. I shall consider some that in my opinion are especially important.

The core of what was just called the explanationist defence of scientific realism is the claim that scientific realism is the best explanation for the impressive predictive accuracy of modern scientific theories and the fruitfulness of the methodology used to obtain them. This claim has an initial plausibility, but to turn it into a defence of scientific realism, two theses will have to be argued for. First, it must be shown that scientific realism is indeed the best explanation for the mentioned successes. If it is, then secondly it must be shown that we may conclude that scientific realism is true or probably true. In other words, it must be shown that the rule of inference called abduction or Inference to the Best Explanation (IBE), according to which explanatory force is a guide to truth, is reliable.

Psillos summarises the standard realist argument for the reliability of IBE as follows:

> [T]he best explanation of the instrumental reliability of scientific methodology is that background theories are relevantly approximately true. These background scientific theories have themselves been typically arrived at by abductive reasoning. Hence, it is reasonable to believe that abductive reasoning is reliable: it tends to generate approximately true theories. (p. 80)

Critics have pointed out that this conclusion follows only if we assume that the fact (if it is a fact) that the approximate truth of the background theories best explains the reliability of scientific methodology makes it reasonable to believe that they are approximately true. But to assume this is to assume the rule of inference the reliability of which is at stake, namely IBE. This has led Fine and Laudan (among others) to the conclusion that the argument is of no significance. According to Psillos, however, this is not necessarily so. Following Braithwaite, Psillos distinguishes between premise-circularity and rule-circularity. An argument is premise-circular if its conclusion is among the premises. A rule-circular argument, on the other hand, is an argument the conclusion of which asserts something about an inferential rule of which the very same argument makes use. As Psillos makes clear, the above argument for IBE is only rule-circular, and such an argument, he argues, need not be viciously circular (even though a premise-circular argument is).

If I understand him correctly, in Psillos' view an argument for the reliability of some rule R essentially involving R is not vicious, because the simple fact that R may be used in any derivation does not guarantee that an argument for the reliability of R can be come by. So, in particular, there might have been no argument for the reliability of IBE even though the use of IBE is sanctioned. In the argument cited above, IBE is needed to arrive at the intermediate conclusion that the background theories are approximately true. But it might have been the case that the approximate truth of those theories is not the best explanation for the success of scientific methodology (p. 83). The argument thereby would no longer be valid.

One may still think that, although its conclusion is not guaranteed to follow, the above argument for IBE is worthless. For, one might say, consider Inference to the Worst Explanation (IWE), a rule of inference that gives us license to infer to the truth of the worst explanation for the data, i.e., the hypothesis that, if true, would make the data more puzzling than any other hypothesis. Suppose IWE were the *modus operandi* in science, and not, as realists claim, IBE. Presumably this would result in very unsuccessful scientific theories. But now consider the following 'justification' of IWE:

> Scientific theories are generally quite unsuccessful
> These theories are arrived at by application of IWE
> ___
> IWE is a reliable rule of inference$_{(\mathrm{IWE})}$

Surely, that IWE is a reliable rule of inference is the worst explanation for what we assume to be the data. So, if IWE were the accepted inferential practice in science, we could, using a rule-circular (but not premise-circular) argument, establish its reliability. But if rule-circular arguments can yield such counter-intuitive conclusions, then it seems there must be something deeply wrong with rule-circularity (even if rule-circular arguments are not viciously circular). (It will be noted that this argument is modeled after Salmon's [1966, pp. 12–17] famous argument against the justification of induction by means of induction itself.)

At this point Psillos' insistence that 'all that is required [in a rule-circular argument] is that one should not have reason to doubt the reliability of the rule'—that there is nothing currently available which can make one distrust the rule (call this principle P)—becomes relevant. After all, it is clear that in my justification of IWE, principle P is not satisfied: There is ample reason to distrust IWE (didn't I say it would presumably lead to unsuccessful theories?). What this so-called justification does show, however, is how crucial P is in Psillos' defence of IBE and, consequently, in his defence of scientific realism. But what is the status of this principle P? Why should we believe it? In particular, should we not require, prior to using some rule in an argument (any argument), that there be reason to believe that the rule is reliable instead of just demanding, as P does, that any doubts about its reliability are absent?

These questions have, in more general terms, attracted a great deal of attention in modern epistemology. Traditionally, it was thought that in order for someone to know (or justifiably believe) that A, the person must at least have good reasons to believe that A; this position is now called internalism. More recently, so-called externalists have argued against this: what is required for knowledge that A is that the belief that A be the result of some reliable belief-forming mechanism. Whether the mechanism is reliable need not be accessible for the person having the belief in order for him or her to know that A, though sophisticated versions of externalism do require that the person not have any reason to doubt either A or the reliability of the mechanism via which s/he came to believe A. Applied to beliefs about rules, these externalist tenets of course yield principle P. Psillos is very explicit about his adherence to externalism, and acknowledges that his defence of realism depends on the tenability of externalism.

It might seem, however, that even if we accept externalism, Psillos' argument for IBE is less than satisfactory. Externalists no doubt will agree with the following passage concerning rule-circularity:

> What is special with rule-circular arguments is what the conclusion says. It asserts that the rule of inference is reliable. But the correctness of this conclusion depends on the rule being reliable, and not on having any reasons to think that the rule is reliable. No less than the conclusion of any first-order ampliative argument, the conclusion of a rule-circular argument will produce a belief, this time about the rule of inference itself. This belief will be justified if the rule is reliable (p. 84).

Note, however, that it is one thing to be justified in one's belief that IBE is reliable, but quite another to be in the position to convince one's opponent that belief in IBE's reliability is justified. And although Psillos is certainly right that for the truth of the conclusion of some argument all that matters (apart from the truth of the argument's premises) is that the rule(s) of inference involved be reliable, the persuasiveness of the argument depends, among other things, on what reasons there are to believe that the rule is reliable. Here externalism on its own cannot help. Suppose someone grants externalism but just does not believe that IBE is reliable (or wants to remain agnostic about its reliability—surely his or her avowal of externalism does not preclude him to have any of those attitudes towards IBE). Then he or she will resist application of IBE in any argument and thus not accept the (correct, in case IBE is reliable) conclusion of Psillos' argument that IBE is reliable, this despite of the fact that he or she, like Psillos, endorses an externalist epistemology. What has Psillos' argument for IBE to offer that could convert this person?

Psillos makes it quite clear that his goal is not to convert the disbeliever in IBE; it is not "to justify [IBE] without any assumptions, or to prove that [IBE is] rationally compelling for any sentient being" (p. 89). Rather Psillos' goal is to give reassurance, from a realist perspective, that IBE is a reliable rule of inference. Or, as one might put it, his goal is to justify IBE 'from within'. Anyone who thinks this is a disappointingly modest goal should realise that it is not at all clear that, when it comes to justifying some fundamental inferential principle, anything more can be asked for. Psillos points to the analogous problem of justifying our deductive practices. We can prove the soundness of our deductive logic, but the proof at the meta-level makes use of exactly the inferential principles the soundness of which (at the object-level) it is trying to establish. This does not make the proof insignificant; it might not have been possible to prove the soundness of deductive logic even though the use of the principles under scrutiny is allowed to that end. Still, it is evident that the soundness proof will not be able to convince someone who refuses to accept deductive logic.

There is one worry concerning Psillos' argument that I do not quite know how to put to rest. The 'justification' of IWE was spurious because it did not respect principle P, we said. But does Psillos' justification of IBE respect P? That is to say, is there really no reason to doubt the reliability of IBE? I think scientific realists are right when they say IBE is commonly used not just in science but also in everyday life. Now, I am sure it has happened to all of us that IBE led us to believe something that we later found out was false. If I have experienced such failure a couple of times, can I still hold that there is no doubt about IBE's reliability? Of course, such failures far from demonstrate that IBE is not reliable: The rule may still mostly lead to true conclusions. Yet they seem to me to be sufficient to make us wonder whether IBE really is reliable. In fact, a natural response to the kind of failures of IBE we sometimes experience seems to be that we try to test IBE, i.e., try to find out the true/false ratio of beliefs obtained by IBE. Such a test might then remove any doubt about IBE but such a test is not part of Psillos' argument for IBE.

In van Fraassen (1989) we also find more principled objections against IBE that at the very least seem to give reason to doubt the reliability of IBE. To mention just one problem van Fraassen raises: if we assume that in general the best explanation for the data is true, then it still only follows that IBE mostly leads to true conclusions if we can be sure that the truth is mostly among the potential explanations we consider. But, van Fraassen asks, given that we have no reason to believe we are privileged in the required way, how can we trust IBE? Psillos (Chapter 9) takes issue with this and other objections and argues that they miss their mark. The counter-arguments against van Fraassen's critique he offers are of two kinds. On the one hand, he tries to show that without IBE, van Fraassen's own antirealist position issues in a blanket skepticism. On the other hand, he tries to meet van Fraassen's arguments directly by arguing that it is not at all implausible to assume we are privileged in the sense required for IBE to be reliable. According to Psillos, we are privileged in that the generation and choice of theories is always guided by background knowledge. But I have some qualms about Psillos' response to van Fraassen.

First, even if it is true that van Fraassen's renunciation of IBE commits him to skepticism, that does not show his arguments against IBE cannot be correct; it would just show that, if they are, both realists and antirealists are in deep trouble. Perhaps in that case we should all become skeptics. Many may find this an unacceptable conclusion, but it should be noted that according to Psillos (p. 215) the skepticism to which van Fraassen is committed if he sticks to his rejection of IBE is inductive skepticism and not Cartesian skepticism. I agree that, if a philosophical position entails Cartesian skepticism, then that counts heavily against it. But it seems to me that the work recently done in formal learning theory (cf. in particular Kelly 1996) at a minimum shows that inductive skepticism is not so obviously absurd.

Secondly, and more importantly, by assuming that we can legitimately speak of a background knowledge privilege, Psillos' positive arguments against van Fraassen's critique of IBE themselves assume IBE. Psillos of course is aware of this and, when presenting his arguments against van Fraassen, refers to his earlier defence of IBE (p. 217). However, that defence was in one important respect left unfinished: It was still to be seen whether the defence of IBE satisfies principle P. And, as Psillos (p. 86) notes, it does not do so if van Fraassen's objections cannot be rebutted. Thus I see a serious problem for Psillos arising from the interdependency of his explanationist defence of scientific realism in Chapter 4 and his counterarguments against van Fraassen's critique of IBE presented in Chapter 9. The problem, spelled out in somewhat more detail, is this: Psillos' argument in Chapter 4 to the effect that we do have background knowledge depends, *inter alia*, on IBE. Whether in arguing for this background knowledge privilege, it is legitimate to use IBE in the way he proposes depends, by Psillos' own lights, on whether principle P is satisfied, i.e., on whether there are reasons to distrust IBE. Now, at least *prima facie* van Fraassen's critique of IBE gives us such reasons. So Psillos should certainly make clear that this critique does not really cast doubt on IBE. However, in doing so Psillos cannot appeal to a background knowledge privilege he can only lay claim to if there are no reasons to doubt IBE. In other words, Psillos cannot in Chapter 9 rely on the conclusion of Chapter 4, for that conclusion had an important proviso, namely that the critique of IBE to be dealt with in Chapter 9 could be met.

To forestall misunderstanding, I do not believe that van Fraassen's critique of IBE shows that this rule cannot be reliable. Nor am I aware of any other arguments to that effect. But, as I said, in my view the foregoing arguments give at least some reason to doubt IBE. Now I take it that principle P is not to be read as saying that we can freely make use of a rule of inference unless it has been demonstrated that the rule is unsound or, in case of an inductive rule, unreliable (if it is, then, as anyone will agree, P is implausibly liberal). So then Psillos will have to say more about why even in the face of the foregoing there is no reason to doubt IBE, or at least no reason to doubt it in the sense meant by principle P.

I should emphasise that nothing of the latter critique is meant to diminish my great admiration for Stathis Psillos' book. In my opinion, it represents not a step but a leap forward in the defence of scientific realism.

Department of Philosophy,
Utrecht University,
Utrecht,
The Netherlands.

By Otávio Bueno

Although the search for a true description of the world is a crucial feature of scientific realism, realists usually grant that we are typically unable warrantedly to assert that a given theory is true. This is particularly the case with theories that deal with unobservable entities or with aspects of the world that are spatio-temporally remote. Theories usually involve idealisations (such as point-masses or ideal gases) and simplifications (to allow, for example, the application of mathematics to the physical domain under consideration). Theories also disregard a number of distorting features, such as the presence of air-resistance in the description of the law of free fall. Furthermore, experimental results often contain errors, and predictions almost never exactly match these results. So it is not surprising that for the scientific realist it is not truth but "truth-likeness [which] is the working notion of truth in science" (p. 276). If exact truth cannot be had, the realist can at least adopt truth-likeness.

But if truth-likeness plays such a role, can the realist provide a sensible account of this notion? Of course, this is a topic that has been in the realist's agenda for a long time. And Psillos spends some time reviewing the difficulties faced by several attempts at providing a formal account of truth-likeness, including Popper's proposal to define verisimilitude in terms of the truth content and the falsity content of a theory, the 'possible worlds' approach, and the 'type-hierarchies' view (pp. 261–275). According to Psillos, none of these accounts works. But the good news for realists is that they don't actually need to provide a formal account of truth-likeness. There is no need to move beyond the intuitive notion of truth-likeness, which Psillos takes to be well enough understood. But what exactly is such an intuitive notion? Here is Psillos's version (following previous works by Weston and Lewis): 'A description D is approximately true of a state S if there is another state S′ such that S and S′ are linked by specific conditions of approximation, and D is true of S′'. As Psillos points out, according to this account, a theoretical law is approximately true of the world if it is true in a world that approximates our world under certain conditions.

But wait: why exactly don't realists need to provide a formal account of truth-likeness? In Psillos's view, because the notion of truth-likeness (as opposed to the notion of truth) is not open to known paradoxes, such as the Liar. As we all know, the existence of such paradoxes led Tarski to formulate a formal account of truth. In the absence of corresponding paradoxes for truth-likeness, the intuitive notion is good enough.

The problem with this response is that, as characterised above, the intuitive notion of truth-likeness is defined in terms of truth. So if truth is open to paradoxes, so is the intuitive notion of truth-likeness. The realist may reply that he or she can always adopt a suitable formal account of truth to avoid the paradoxes. The difficulty with this reply is that if Psillos defines truth-likeness in terms of a formal account of truth, he can no longer claim to have an intuitive notion of truth-likeness. With truth-likeness defined in terms of formal truth, truth-likeness becomes formal too.

More importantly, the realist also needs to establish the connection between truth and truth-likeness (and one would expect something more than a simple definitional connection). It is not by chance that some realists go on to assert that truth is a limiting case of truth-likeness (see p. 273, where the claim is made that truth is a limiting case of versimilitude). But it is unclear how the above intuitive account of truth-likeness can be used to maintain such a connection between truth and truth-likeness. How can truth be a limiting case of truth-likeness if in order to define truth-likeness the notion of truth is presupposed?

Perhaps the realist could reply that truth is not, after all, a limiting case of truth-likeness. But this seems to leave the realist with a notion of truth-likeness that isn't of much use for realism. Realists are, of course, ultimately concerned with truth. Truth-likeness is, as Psillos points out, just the "working notion of truth in science" (p. 276). The adoption of truth-likeness can be seen as a pragmatic expedient given the messiness of scientific practice. But I take it that the idea is to eventually get to truth—via truth-likeness. But how can this be done?

Well, it is not clear at all that it can be done. The difficulty here is that, according to the intuitive account, judgments about truth-likeness are thoroughly context-dependent. After all, the 'specific conditions of approximation' that are used in the definition of truth-likeness change from one context to another. But truth is not context dependent—especially for the realist. To allow truth to be context dependent is, of course, to open the door to all sorts of relativisms that are anathema for a realist view. So a considerable gap between truth and truth-likeness needs to be bridged. But it is unclear how the intuitive notion of truth-likeness can be used to bridge this gap, given that it crucially depends on the thoroughly context dependent 'specific conditions of approximation'. The intuitive notion, if anything, seems to highlight the gap.

This seems to leave the scientific realist in an unstable situation: there is the need for truth-likeness in the realist view, given the difficulties to assert the truth of a theory (due to the presence of idealisations, simplifications and so on). However, neither the formal account of truth-likeness nor the intuitive view seem to work. Given this scenario, perhaps the move to something weaker than truth, such as empirical adequacy, may not look so bad after all.

In order for Psillos to provide an account of the maturity of a scientific theory, he needs to articulate an account of the indispensability of a hypothesis to make a successful prediction. The maturity issue is, of course, crucial for the realist, given Psillos's own formulation of scientific realism: according to the scientific realist, only mature scientific theories are expected to have terms that refer, and these are the theories that are expected to be truth-like (p. xix). But under what conditions do we say that a hypothesis indispensably contributes to the generation of a successful prediction? According to Psillos, the predictive indispensability of a hypothesis can be characterised thus:

> Suppose that H together with another set of hypotheses H′ (and some auxiliaries A) entail a prediction P. H indispensably contributes to the generation of P if H′ and A alone cannot yield P and no other available hypothesis H* which is consistent with H′ and A can replace H without loss in the relevant derivation of P (p. 110).

As Psillos correctly notes, there is a sense in which no theoretical claim is indispensable to the generation of a given prediction: we can simply adopt a Craig-transform of the theory under consideration or we can "'cook up' a hypothesis H* by writing P into it" (p. 110). To avoid these possibilities, Psillos advances some epistemic constraints that any putative hypothesis is required to satisfy. The constraints include requirements that a theory be independently motivated, non-ad hoc, potentially explanatory and so on. With these constraints, Psillos can then conclude: "it is not certain at all that a suitable replacement can always be found" (p. 110). I agree.

The problem, however, is that with those epistemic constraints, it is not clear that any suitable replacement can ever be found. It might be said that this is exactly what Psillos intended. After all, the point of having an account of predictive indispensability is precisely to single out one—and only one—crucial hypothesis without which the successful prediction can't made. And it is clear why the scientific realist needs an account as strong as this: without such a tight connection between the relevant hypothesis and the successful prediction, scientific realists would have a hard time avoiding Laudan's well-known counterexamples of theories whose predictions seem to be successful but whose terms don't refer (see Laudan [1981] and Laudan [1996]). To deflate Laudan's counterexamples, Psillos strategy is to single out the class of predictively indispensable hypotheses, and claim that those are the hypotheses the scientific realist should be realist about. Given that such hypotheses are indispensable to the generation of the relevant predictions, we cannot abandon the hypotheses without losing the predictions.

But in trying to avoid Laudan's counterexamples, Psillos's account seems to have moved too far. For how can the scientific realist establish that a particular hypothesis indispensably contributes to the generation of a given prediction? In order to establish that, not only does one need to establish that there is a hypothesis H such that, together with other auxiliary hypotheses H′ and background assumptions A, it predicts P—which can be done—but one also needs to establish that no other 'available' hypothesis H* predicts P (together with H′ and A). In other words, what the realist needs to establish is that H is the one and only hypothesis that generates the prediction of P (given H′ and A). But how can the realist ever be able to establish that?

The problem here emerges from the requirement that there is no available hypothesis H*. What does the realist mean by 'available'? If by 'available' it is meant 'at the historical moment in which H was originally entertained', then Psillos's account of predictive indispensability lacks the force to overcome Laudan's criticism. For there will be hypotheses that were predictively indispensable at one time (given that there were no available alternatives to them when the hypotheses were first formulated), but which turned out not to be indispensable at a later time (with the formulation of a new hypothesis that generated the relevant prediction). With this account, the scientific realist will be basically endorsing Laudan's criticism!

However, if by 'available' the realist means 'conceptually available', rather than historically so, in the sense that the alternative hypothesis H* may not have actually been available at the moment in which H was first formulated, but could be entertained by some scientific community in the future, then Psillos's account becomes way too strong, even for the realist. For how could the realist assert that a given hypothesis is indeed predictively indispensable? He or she would need to show that there is no conceptually available hypothesis H* that also generates the prediction P. But unless the hypothesis H* is inconsistent (a possibility that the scientific realist won't take seriously anyway), a whole range of hypotheses H* are conceptually available (and they even satisfy Psillos's epistemic constraints). In other words, the difficulty here is that, at any particular moment of time, the realist will never be able to establish that there are no conceptually available hypotheses H* that entail P—and so the realist won't be able to establish that H indispensably contributes to the generation of P. For even if the scientific community may not be able to conceive of H* in a particular moment, this doesn't entail that it won't be able to conceive of it in the future.

However, perhaps the interpretation of 'available hypothesis' as sheer conceptual availability is enough for the scientific realist's needs. For it allows him or her to deal with past cases of predictive indispensability of scientific theories. Although the scientific realist is unable to claim that a given hypothesis is predictively indispensable, he or she can at least claim that a particular hypothesis is not indispensable. For if a suitable replacement H* has ever been found for a given hypothesis H, that's all that the realist needs to claim that H wasn't indispensable after all.

But this leaves the scientific realist in an awkward position. He or she can never be warranted in asserting that a given hypothesis is predictively indispensable, and so he or she can never be warranted in asserting that a given theory is indeed mature. At best the realist can claim that a given hypothesis isn't predictively indispensable, and so that a given theory isn't mature after all. What this entails is that the scientific realist is unable to assert that he or she is realist about a given hypothesis H, since in order to make this assertion, the realist will need to establish that H is predictively indispensable. At best, the realist can claim that he or she is not realist about a certain hypothesis H′—supposing that H′ is not predictively indispensable, and so it's only an 'idle' component. To say the least, it is unfortunate not to be able to assert one's own position.

Furthermore, this generates a problem for the scientific realist. The reason why the realist needs an account of predictive indispensability is to avoid having to provide an account of success for 'idle' theoretical components; only the indispensable components in a theory need explanation. However, without being able to establish that a given hypothesis is indeed predictively indispensable—without being able to assert that a given hypothesis contributes essentially to the predictive success of a theory—the realist cannot claim to have provided an explanation of the success of science. For the realist cannot warrantedly assert that these are the indispensable components responsible for the theory's success—not, at least, with the account of predictive indispensability provided by Psillos.

It is unfortunate that Psillos only discusses Worrall's epistemic version of structural realism (Worrall [1989]), neglecting Ladyman's ontic version (Ladyman [1998]) to a two-line footnote (p. 309, note 5). For the ontic version is not open to Psillos's charges against structural realism (see also French [1999] and French and Ladyman [2001]).

Psillos's main complaint against structural realism is that it cannot explain the continuity of structure in scientific change without simply falling back into scientific realism. The idea is that, in isolation from a number of theoretical assumptions, a mathematical structure cannot explain anything about the world. So a commitment to structure is not enough to explain the retention of mathematical structures in theory change. According to Psillos, the same cannot be said about scientific realism. After all:

> Scientific realists can explain the fact that mathematical equations have been retained in theory change by saying that they form an integral part of the well-supported and (approximately) true theoretical content of theories. But they would deny that all of what is retained is empirical content and (uninterpreted) mathematical equations. Not only is some theoretical content also retained, but scientists now have good reason to believe that the content of current theories—what they predicate of the world—is better supported by evidence, and, hence, more likely to be true (p. 147).

Moreover, Psillos continues:

> The fundamental insight Worrall has, i.e. that the predictive success of a theory points to the theory's being correct in some of its claims about the unobservable world, cannot be best served by a distinction along the lines of structure (or mathematical equations) versus nature (or theoretical content)... The best place to draw the relevant line is between essentially contributing theoretical components and 'idle' ones (p. 155).

It is difficult to see how Psillos can maintain that the scientific realist is able to explain the retention of mathematical structure in theory change. For, according to Psillos (see the second quotation above), to explain such structural preservation requires a distinction between "essentially contributing theoretical components and 'idle' ones". But, as argued above, Psillos attempt to characterise the latter distinction fails. And so it's not clear that the scientific realist is actually able to provide the intended explanation.

Of course, the distinction between essentially contributing theoretical components and 'idle' ones has no role in structural realism, and it is not a distinction that the structural realist is committed to. The explanation of the retention of mathematical structures in theory change is accomplished by a claim about the adequacy of such structures to represent the relevant structures of the world. This is one of the reasons that motivated the move to the ontic version of structural realism in the first place (see Ladyman [1998], French [1999], and French and Ladyman [2001]). In this respect, if nothing else, the structural realist seems to be in a better position than the scientific realist to accommodate scientific change.

Even if we grant that structural realism has its problems, I don't think that Psillos's case against this position succeeds. Structural realism, particularly in Ladyman's ontic version, still seems to remain the best hope for the realist.

Department of Philosophy,
California State University,
Fresno, CA 93740-8024, USA.

Author's Response

By Stathis Psillos

It's a privilege to have one's book discussed so thoroughly and carefully as the reviewers do with my book. So, I thank them all for their incisive and thought-provoking criticism. Lack of space doesn't allow me to discuss their points in the detail they deserve. So, I'll restrict my attention to two central issues that crop up in the reviews: the role of the 'No Miracles Argument' (NMA) in the defence of realism and the prospects for Structural Realism.

Lipton poses a strong challenge: show that NMA does some genuine extra work for realism over and above the work already done by first-order instances of Inference to the Best Explanation (IBE) that scientists use in order to form their theories about the unobservable world. Lipton and I both agree that circular vindications of inferential methods can be legitimate. His challenge is peculiar to an NMA-like vindication of IBE: unlike legitimate inductive assessments of inductive methods, an NMA-like assessment of the reliability of IBE introduces "no new evidence" for its reliability. Let me remind the reader that I take NMA to consist in two parts. The first part is that we should accept as relevant (approximately true) the theories that are implicated in the (best) explanation of the instrumental reliability of first-order scientific methodology. The second part is that since, typically, these theories are arrived at by IBE, IBE is reliable. Both parts are necessary for my version of NMA. Note that the second part has excess content over the first. The first part says nothing yet about the reliability of IBE. On the contrary, the second part issues a general statement about a mode of reasoning which, being contingent, implies that there must be a feature of the world that answers to IBE's reliability. That the world is such that—as a contingent matter of fact—IBE tends to yield (approximately) true theories is a new general claim about the world which is not entailed by the (scientists') first-order IBEs. In support of this, note that the general claim (that IBE is reliable), if true, would support relevant counterfactuals which no first-order claim could support.

But perhaps all this doesn't meet Lipton's worry that there is no new evidence for the second part of NMA—viz, that IBE is reliable—which is not already evidence for the "first-order scientific case for [the] truth [of theories]". A lot seems to depend on how to understand the requirement of "new evidence". But two things seem relevant here. First, the actual track-record of successful applications of IBE does offer genuine evidence for the reliability of IBE. In particular, successful novel predictions issued by first-order theories arrived at by IBE do lend extra credence to the claim that IBE is reliable. Second, the reliability of IBE offers new evidence (of a sort) for the truth of first-order scientific theories. Suppose theories T_1, $\ldots T_n$ are accepted as true on the basis of the relevant evidence. Suppose also that one grants that the fact that these theories were arrived at by a method (viz., IBE) offers no new evidence for the claim that they are (probably) true which was not already there by considering the (first-order) evidence for them. Yet, it seems to offer some new evidence for the claim that a fresh theory T_{n+1}, which is arrived at by IBE (perhaps, a theory from a totally different domain), is (probably) true. The successes of $T_1, \ldots T_n$ and the fact that they were arrived at by IBE, supports, via NMA, the view that IBE is reliable, and this works in addition to the first-order evidence for T_{n+1} to make T_{n+1} more credible. Generally, the fact that a theory has been arrived at by a reliable method will have extra probative force: it will add to our confidence in its (probable) truth, for its truth will also be supported (indirectly) by all the (first-order) evidence that has led scientists to accept the method as reliable.

There is one more reason why NMA has genuine excess content over and above first-order IBEs. The latter are diverse and disparate (e.g., they might admit the form of common-cause arguments, or of arguments from unification etc.). If NMA is correct, then it says something about the common deep inferential structure of the several instances of explanatory reasoning and suggests that a host of possibly disparate instances of this success-to-truth-via-explanatory-considerations mode is reliable.

Douven questions whether Principle P (which I take it to mean that a doxastic practice is innocent until good reasons show it to be guilty) is satisfied in the case of IBE. He takes van Fraassen's argument from the bad lot (henceforth, ABL) to pose such doubts for the reliability of IBE. I disagree. The *prima facie* doubt that ABL is supposed to cast on IBE should be this: given the plausible principle 'no truth in, no truth out', one can doubt that, so to speak, there is truth in. That is, one can doubt that IBE operates within an environment of true or truthlike theories. But this is not a doubt about the reliability of IBE unless it is accompanied by an argument which shows that IBE is unlikely to operate within an environment of truthlike theories. The argument from the bad lot doesn't establish this. All ABL does is pose a (legitimate) challenge to realists: given the cogent principle 'no truth in, no truth out', offer some good reasons that IBE does operate in an environment of true or truthlike theories.

Douven's construal of ABL seems to me unwarranted. He asks: "given that we have no reason to believe we are privileged in the required way [i.e., that there is truth in], how can we trust IBE?". For one, ABL does not show that we have no reason to believe that we are privileged in the required way. For another, the realists can try to meet the challenge in the following way. On the one hand, they can argue—negatively—that we are given no reason to doubt that scientists are so privileged. On the other hand, they can argue—positively—that, on the basis of NMA, we can explain how scientists can be so privileged: if NMA is cogent, then we have reasons to trust IBE. This last move does depend on Principle P, as Douven correctly stresses. But this dependence is legitimate. There is no reason before we appeal to Principle P to discharge any doubts that ABL might pose on the application of this Principle to the case of IBE, since ABL does not cast doubt on the reliability of IBE.

Both Redhead and Bueno take me to task on the issue of Structural Realism (SR). They raise a number of interesting points and I am certainly willing to learn from their criticism. Bueno, for instance, rightly observes that I neglected the ontic version of Structural Realism. I have tried to rectify this fault (forthcoming).

My reaction to epistemic SR, on which I still insist, was that this view is either trivially true or belies structuralism. The route I chose to show this was (following Maxwell) to associate SR with a Ramsey-sentence understanding of theories and to show that Newman's challenge to Russellian structuralism applies to the Ramsey-sentence case equally forcefully. Redhead rightly objects that if we don't understand Worrall's SR as tied to a Ramsey-sentence style structuralism, then my objection (partly) fails. He admits that I was right to emphasise that "the reference of $\bar{\mathrm{R}}$ must be picked out in non-structural terms". I took this admission to belie structuralism, but Redhead thinks that it's consistent with Worrall's SR. His claim "is merely that $\bar{R}$ is hypothesised in some explanatory theoretical context so it exists as an ontological posit, but all that we have epistemic warrant for is the second-order structure S". I doubt that this claim is cogent. An example will illustrate my worry. Take the relation $\bar{R}$ ('has more mass than') which is instantiated by the particles proton and electron. Suppose that we posit this (definite, à la Russell) relation $\bar{R}$ in some explanatory theoretical context. This relation has a formal abstract (second order) structure. In particular, it has the structure (it belongs to the isomorphism class) of a relation which is irreflexive, asymmetric and transitive—call it S. On Redhead's view, it follows that we are only warranted in our claims about this abstract structure S. Now, even if we were to accept that the Newman challenge does not apply any more (though I think it still does), I cannot see how we can have epistemic warrant to believe that the structure is S which is not parasitic on having epistemic warrant to believe that $\bar{R}$ is the required (definite) relation. It is the knowledge of (or commitment to) the definite relation $\bar{R}$ ('has more mass than') that issues the warrant to believe that its structure is S and hence issues the warrant that the relevant domain has structure S. More generally, the epistemic warrant moves from the concrete (interpreted) structure that we hypothesise in a certain theoretical context to the relevant abstract structure. Obviously, once this is done one can reverse the epistemic order. Since belief in S is weaker than belief in $\bar{R}$, one can retreat to an epistemically safer belief. After all, it is the case that $\mathrm{prob}(\bar{R}) \leqslant \mathrm{prob}(S)$. But although this move is legitimate, by making the epistemic version of SR to recommend just this move we make it lose most of its original appeal. It's no longer the claim that only structure can be known (epistemically accessed). The above probabilistic relation says nothing about how high $\mathrm{prob}(\bar{R})$ is and hence it says nothing about whether we are entitled to believe in $\bar{R}$ or not. Where SR aimed to issue in a principled epistemic constraint, it now issues a mere word of caution.

Bueno argues that my alleged inability to show which constituents of a theory are idle and which are essentially contributing to its successes makes Structural Realism be "in a better position than the scientific realist to accommodate scientific change". Briefly put, I characterised as indispensably contributing to the successes of theories those hypotheses which cannot be replaced by other available hypothesis H* in the generation of a successful prediction. Bueno's challenge is this: if I mean no other "actually" available hypothesis H*, then my account is too weak, for there may be such hypotheses H* in the future. If, on the other hand, I mean no other "conceptually" available hypotheses H*, then my account is too strong, since there is no way to exclude the conceptual possibility that such hypotheses H* are available. By way of reply, I note that I stick to actuality. I don't want to exclude the conceptual possibility that an IH hypothesis might prove dispensable. It might and yet it might not. But if we were to take seriously the thought that no hypothesis H can be said to be indispensably contributing to successes unless it is conceptually impossible that there is an H* which might render H dispensable, we would require that empirical science give way to a priori theorising. So, what if we restrict the scope of IH to those hypotheses that don't have actual rivals H*? I don't share Bueno's confidence that "there will be hypotheses that were predictively indispensable at one time . . . but which turned out not to be indispensable at a later time". This sounds more like an open empirical issue to be judged by the relevant evidence. In any case, even if a hypothesis H (which is now taken to contribute indispensably to some successes) gets replaced by a future H*, it may still be the case that H can be seen as approximately true in the light of H*. I think this is quite likely, if the intuition behind my original suggestion about IH is correct. For if success—and especially novel success—of a theory points to the theory's having truthlike constituents, those constituents that have actually led to successes won't be characteristically false. Hence, they may well either be retained as they are or else be seen as approximately true in the light of a replacing theory. If the worry is how we can tell whether currently indispensable hypotheses will remain so (even approximately) in the future, I refer the reader to Sklar (2000) for a sketch of a "methodology of theoretical anticipation".

Department of Philosophy & History of Science,
University of Athens,
15771 Athens, Greece.

References

French, S. (1999). "Models and Mathematics in Physics: The Role of Group Theory", in J. Butterfield and C. Pagonis (eds), *From Physics to Philosophy*. Cambridge: Cambridge University Press, 187–207.

French, S., and Ladyman, J. (2001). "Remodelling Structural Realism: Quantum Physics and the Metaphysics of Structure", forthcoming in *Synthese*.

Howson, C. and Urbach, P. (1993). *Scientific Reasoning: The Bayesian Approach*, 2nd edn., Chicago: Open Court.

Kelly, K. (1996). *The Logic of Reliable Inquiry*. Oxford: Oxford University Press.

Ladyman, J. (1998). "What is Structural Realism?", *Studies in History and Philosophy of Science*, 29: 409–424.

Laudan, L. (1981). "A Confutation of Convergent Realism", *Philosophy of Science,* 48: 19–49.

Laudan, L. (1996). *Beyond Positivism and Relativism*. Boulder: Westview Press.

Newman, M. H. A. (1928). "Mr Russell's 'Causal Theory of Perception'", *Mind,* 37: 137–148

Psillos, S. (forthcoming). "Is Structural Realism Possible?" *Philosophy of Science*.

Salmon, W. (1966). *The Foundations of Scientific Inference*. Pittsburgh: University of Pittsburgh Press.

Sklar, L. (2000). *Theory and Truth*. Oxford: Oxford University Press.

Worrall, J. (1989). "Structural Realism: The Best of Both Worlds?", *Dialectica,* 43: 99–124.

