Studies in History and Philosophy of Science 44 (2013) 563–570

Contents lists available at ScienceDirect

Studies in History and Philosophy of Science

journal homepage: www.elsevier.com/locate/shpsa

Normativity, the base-rate fallacy, and some problems for retail realism

Paul Dicken

Munich Center for Mathematical Philosophy, Ludwig-Maximilians-Universität, D-80539 München, Deutschland, Germany
School of Humanities and Languages, University of New South Wales, Sydney, NSW 2052, Australia

ARTICLE INFO

Article history:
Received 10 June 2013
Received in revised form 17 September 2013

Keywords:
Scientific realism
No-miracles argument
Base-rate fallacy

ABSTRACT

Recent literature in the scientific realism debate has been concerned with a particular species of statistical fallacy concerning base-rates, and the worry that no matter how predictively successful our contemporary scientific theories may be, this will tell us absolutely nothing about the likelihood of their truth if our overall sample space contains enough empirically adequate theories that are nevertheless false. In response, both realists and anti-realists have switched their focus from general arguments concerning the reliability and historical track-records of our scientific methodology, to a series of specific arguments and case-studies concerning our reasons to believe individual scientific theories to be true. Such a development however sits in tension with the usual understanding of the scientific realism debate as offering a second-order assessment of our first-order scientific practices, and threatens to undermine the possibility of a distinctive philosophical debate over the approximate truth of our scientific theories. I illustrate this concern with three recent attempts to offer a more localised understanding of the scientific realism debate—due to Stathis Psillos, Juha Saatsi, and Kyle Stanford—and argue that none of these alternatives offer a satisfactory response to the problem.

When citing this paper, please use the full journal title *Studies in History and Philosophy of Science*

1. Introduction

Recent literature in the scientific realism debate has been particularly concerned with a species of statistical fallacy that appears to undermine both realist and anti-realist arguments regarding the approximate truth of our scientific theories. Howson (2000, pp. 52–54), for example, argues that while the likelihood of a true scientific theory generating successful predictions is of course extremely high, one cannot thereby infer the corresponding probability of a successful scientific theory being true without taking into account the base-rate probability of any arbitrary scientific theory being true—which is of course precisely what the scientific realist is attempting to establish. Similarly, Lewis (2001) complains that while the history of science may furnish us with numerous examples of successful scientific theories that have eventually been proved to be false, one may not thereby infer the probability of our current scientific theories going astray without knowing the underlying probability of a false scientific theory generating successful predictions—which is of course precisely what the sceptic wishes to establish. Magnus and Calendar (2004) go so far as to suggest that this widespread tendency to ignore the base-rates explains the sense of intractability that pervades recent discussion in the scientific realism debate: that while further historical examples and case-studies, or more demanding and selective criteria of success, may raise or lower the likelihoods in question, at the end of the day all such considerations will be simply swamped by the underlying probability of any arbitrary scientific theory being true—something over which realists and anti-realists can only trade intuitions.

The result has been a reconceived focus for the scientific realism debate. In the view of Magnus and Callender (ibid.: 333–336) for example, we should abandon those traditionally *wholesale* arguments for or against scientific realism that rely upon sweeping statistical claims in favour of a series of *retail* arguments targeting more specific cases. Psillos (2009, pp. 65–66) argues in particular that the scientific realist should be concerned with the likelihood of an *individual* scientific theory being true given its predictive success, rather than the abstract relationship between success and

E-mail address: email@pauldicken.com

0039-3681/$ - see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.shpsa.2013.09.005

truth in general. Pursuing a similar line of thought, Saatsi (2010) suggests that we should prefer those arguments for realism that are concerned with the distinctive *content* of each individual inference in question, rather than those that concern the generic *form* of many such inferences; and Stanford (2006) has recently defended a version of instrumentalism that is based upon the epistemological pedigree of particular scientific theories, rather than a general distinction—between, say, the observable and the unobservable—that cuts across our scientific theories without discrimination.

This narrowing of focus has brought with it a renewed emphasis upon detailed case-studies in the history of science that has undoubtedly enriched the debate. Yet it also raises a deeper concern, since as the argument over scientific realism becomes increasingly tied to the specifics of individual scientific theories, it becomes increasingly difficult to see what the distinctively *philosophical* contribution to that debate might be: certainly, one would no longer be able to conclude anything from the general evaluation of the patterns of reasoning exemplified in scientific practice if one also considers such reasoning to be fundamentally context-specific. The scientific realism debate threatens therefore to become just another aspect of the first-order scientific deliberations with which it was supposedly concerned—a conclusion that some critics may welcome, but one that is at least in tension with the explicit intentions of those philosophers who take themselves to be *reformulating* the scientific realism debate, rather than abandoning it altogether.

I illustrate this situation by first outlining the no-miracles argument for scientific realism and how it is understood to suffer from the so-called base-rate fallacy (Section 2), before turning to the various ways in which realists and anti-realists have attempted to restrict the scope of their reasoning to individual scientific theories in response (Section 3). Essentially, my argument is that there is a tension between, on the one hand, targeting individual theories and specific instances of our scientific methodology and, on the other hand, attempting to reason about these theories and inferences in general. I argue that Psillos' reformulation of the scientific realism debate as an attempt to balance our competing first-order and second-order evidence (Section 3.1), Saatsi's distinction between form-driven and content-driven arguments for scientific realism (Section 3.2), and Stanford's selective anti-realism (Section 4), all fail to adequately distinguish a distinctively philosophical dimension to the question concerning the approximate truth of our scientific theories. The paper concludes with some observations regarding how the scientific realism debate's insistence upon an often poorly articulated naturalistic methodology appears to lie at the root of these difficulties (Section 5).

2. The no-miracles argument and the base-rate fallacy

The central argument for scientific realism—the so-called no-miracles argument—is essentially an inference to the best explanation, the rough idea being that we should believe our contemporary scientific theories to be approximately true as this offers the best explanation for their wide-spread predictive success (e.g. Boyd, 1980; Putnam, 1975). There is of course by now a vast literature on the precise details of this argument, concerning such issues as to whether or not it really should be understood as an abductive inference (Musgrave, 1988); and if it is, with whether or not the truth of our scientific theories really is the best explanation for their predictive success (van Fraassen, 1980, pp. 39–40), and why we should suppose our best explanations tell us anything about the world anyway (van Fraassen, 1989, pp. 142–149). More generally however, the worry has been that such reasoning simply begs the question, since it was concerns over the reliability of inference to the best explanation in our first-order scientific practices that motivated the defence of scientific realism in the first place (Fine, 1984).

Nevertheless, such reasoning may still be able to make a positive epistemological contribution to the scientific realist's position. As Psillos (1999, pp. 78–81; see also his 2009, pp. 49–52) understands the argument, the point of the no-miracles argument is not so much to *justify* our confidence in the approximate truth of our scientific theories against the challenges of the committed sceptic, but rather to help *explain* the predictive success of those theories for those who already enjoy a realist disposition.[1] We can think of the no-miracles argument therefore as providing additional reassurance for the scientific realist who, having concluded that a scientific theory is probably true on the basis of its predictive success and other theoretical virtues, may nevertheless entertain the meta-theoretical worry as to why we should trust our scientific methodology in this respect. The answer offered by the no-miracles argument is that *in general* our scientific reasoning is reliable; and while this inference will itself be an inference to the best explanation, and therefore part of the very scientific methodology under dispute, it does at least speak to the overall consistency of the scientific realist's position—a non-trivial result, since it could easily be the case that one's commitment to our first-order scientific practices actually sat in tension with our ability to offer a second-order philosophical argument regarding the reliability of those practices.[2]

More specifically, this *explanationist* defence of scientific realism proceeds via a two-stage process. We begin with the second-order (philosophical) inference: that in general we have good reasons to believe our scientific theories to be true, since this is the best explanation for their predictive success. This is the starting realist intuition, and the purpose of the no-miracles argument is then seen to be to provide further (internally consistent) justification for this intuition. The second step is then to note that the scientific theories we believe to be approximately true were themselves the product of a first-order (scientific) inference to the best explanation—and since *ex hypothesi* we have assumed these scientific theories to be approximately true, we therefore have reasons to believe these first-order inferences to the best explanation to be reliable. So we use an inference to the best explanation to justify inference to the best explanation; but they are at least different inferences to the best explanation, since it is the second-order *philosophical* argument that is being used to justify the reliability of the first-order *scientific* inferences. And this then completes the justificatory circle, for if our first-order scientific methodology is indeed reliable, then we do have good reasons to suppose that in general the scientific theories they produce will be approximately true, and the meta-theoretical doubt shown to be inconsistent with our starting intuition.

The value of such reasoning then lies in demonstrating the overall coherence of the scientific realist's position—the philosophical argument supports the scientific inference, and the scientific inference supports the philosophical argument. Such reasoning however proceeds at an extremely general level of abstraction, and as such is vulnerable to a particular kind of statistical fallacy. Let us grant for the sake of argument that the scientific realist is right, and that we do have good reasons to suppose that our

[1] In a recent paper, Psillos (2011b) qualifies his early claim that the no-miracles argument provides the "ultimate" argument for scientific realism, since it presupposes certain broad methodological assumptions regarding e.g. the desirability of explanation, which are in fact *constitutive* of scientific realism. Nevertheless, he still maintains that the no-miracles argument provides support for scientific realism, and can be used as a vindication for abductive reasoning.

[2] Contrast this with at least one way of understanding the pessimistic meta-induction, as an attempt to argue for the *unreliability* of our predominantly inductive scientific methodology on the basis of an induction upon its historical track-record; c.f. Lipton (2000).

contemporary scientific theories are true, as this provides the best explanation for their predictive success. Let us also grant that these scientific theories were themselves the product of various first-order inferences to the best explanation. The problem however is that these two assumptions alone are not sufficient to help explain the predictive success of our scientific theories—nor *a fortiori* to help demonstrate the internal coherence of the scientific realist's position—since they are in fact consistent with our scientific methodology being extremely *unreliable*. Suppose for example that out of the all the scientific theories under consideration, an extremely large number of them are predictively successful, but only a tiny fraction of these are actually true: we could still have good reasons to believe that our currently accepted scientific theories are true (this being a tiny subset of all the scientific theories under consideration), but without any of this changing the underlying fact that in general, the probability of our scientific methodology generating a true scientific theory is extremely low.

More precisely, the explanationist defence of scientific realism faces a problem regarding the *base-rates* of any arbitrary scientific theory being true. The issue is usually illustrated with a medical example, where we are attempting to identify the prevalence of some disease in a population by testing for its symptoms, but where we could just as easily think of ourselves as attempting to identify the prevalence of truth amongst a population of scientific theories by testing for their predictive success. Suppose that we have such a test with e.g. 95% reliability—meaning here that it correctly identifies the disease in 95% of the cases where it is present, and misdiagnoses the disease in 5% of the cases where it is not. We perform the test on a randomly selected member of the population and get a positive result; what is the probability that the disease is actually present? The answer of course is that given the above information, we simply don't know, since it will depend crucially upon the underlying frequency of the disease in the population. Suppose that our population consists of 100,000 people, but that the disease is only present in 0.1% of cases. Our test will then correctly identify $(0.95 \times 100) = 95$ of the infected population, and will misdiagnose $(0.05 \times 99{,}900) = 4996$ of the uninfected population; the actual probability of our arbitrary patient actually having the disease given a positive test result is therefore given by the ration of correct identification to total identification, which in this case will be less than 2%.[3]

The point of course then is that precisely the same concern holds for the explanationist defence of scientific realism: that while we can happily grant that the likelihood of a true scientific theory generating successful predictions will be extremely high, the corresponding and far more important likelihood of a predictively successful theory turning out to be true will be largely determined by the underlying probability of any arbitrary scientific theory in our overall sample space being true; and the problem here is that this base-rate probability is not something that the scientific realist will be in a position to specify independently of having already endorsed the conclusion of the no-miracles argument (Howson, 2000, pp. 52–54). So the initial threat of circularity remains, even given the more limited motivations and goals of the explanationist defence of scientific realism—the scientific realist can no longer justify his initial philosophical argument on the basis of the reliability of our first-order inferences to the best explanation, since this is in fact precisely what his intended argument attempts to establish; and if that is the case, then the no-miracles argument adds absolutely nothing to the initial philosophical argument, not even a demonstration of its overall consistency. To put the point even more succinctly, any evidence the scientific realist can offer for the truth of our current scientific theories will be swamped by the background probability of an arbitrary scientific theory being true; therefore, in order for the no-miracles argument to make any positive justificatory contribution to scientific realism, one must first assume that the base-rate likelihood of a predictively successful theory being true is actually quite high; but to assume *that* is just to assume what the scientific realist was attempting to establish all along.

The solution has been to distinguish between the no-miracles argument as a piece of reasoning that concerns any arbitrary scientific theory, and the no-miracles argument as applied to specific scientific theories; and to maintain that while the former does indeed depend upon our overall sample space of successful scientific theories, the latter only depends upon features particular to the scientific theory in question—as Psillos notes, "the approximate truth of each and every theory will *not* be affected by the number (or the presence) of other theories... approximate truth, after all, is a relation between the theory and its domain" (2009, p. 65). But now a second worry begins to appear: as it was originally presented, the no-miracles argument was an attempt to further justify the scientific realist's position (if only in terms of demonstrating its overall internal consistency) by showing that the initial philosophical argument was an instance of a *general* inferential pattern that our first-order scientific practice shows to be reliable; but if we are now to attend to the details of specific scientific theories, it is no longer clear that there will be a general inferential pattern that our first-order scientific practice can be said to exemplify.

Resisting the base-rate fallacy therefore raises important questions as to what exactly the scientific realism debate is attempting to achieve. The philosopher of science cannot be simply concerned with the first-order reasoning by which scientists come to accept one specific scientific theory over another, for that would be just to do more science—a task best left to the experts. Or more precisely, while there are undoubtedly interesting philosophical questions to be raised about our first-order scientific practices—measures of confirmation, accounts of explanatory power, clarification and analysis of specific theoretical concepts—none of this was what the scientific realism debate was about, the supposedly *philosophical* question as to whether or not we should believe those scientific theories already selected for by the scientific community.

3. Some problems for retail realism

In their original paper, Magnus and Calendar (2004) distinguish between *wholesale* arguments for scientific realism that proceed along sweeping statistical lines and are therefore vulnerable to worries relating to the base-rate fallacy, and *retail* arguments for scientific realism that are concerned with "specific kinds of things such as neutrinos" (ibid.: 321) and are consequently immune to such objections. The challenge however is to specify exactly what a retail argument for scientific realism amounts to, and in particular, to show that it manages to make a genuine philosophical contribution to the debate without either collapsing into our first-order scientific reasoning, nor operating at such a level of generality that we must again contend with the underlying base-rate probabilities. Magnus and Callender do not themselves have a great deal to say on this issue, noting merely that retail reasoning in the philosophy of science is to answer questions about, for example, the reality of atoms "by referring to the same evidence scientists use to support the atomic hypothesis" (ibid.)—yet without understanding such evidence as part of a more general claim regarding the reliability of our scientific methods. Yet without further clarification, it is hard to see how this amounts to anything more than simply repeating the same first-order

[3] If the base-rate probability of the disease in our population is exactly 50%, then the probability of having the disease given a positive result will be 95%.

reasoning that scientists already offer for accepting the theory in question; and the philosophical issue as to whether or not we have good reasons to suppose that our accepted scientific theories are therefore true is simply left unaddressed.

We can put the general worry another way. A number of philosophers have sought to provide extremely specific arguments for localised forms of scientific realism through the careful reconstruction of particular episodes in the history of science: for example, both Salmon (1984) and Achinstein (2001)—and indeed, many others—have argued that Perrin's work on the experimental confirmation of Avogadro's number constitutes not only reasons for accepting the atomic hypothesis of matter, but moreover, reasons for believing that atoms really exist. Such reconstructions therefore offer to provide a paradigmatic example of a retail argument for scientific realism, and are presumably exactly the sort of thing that Magnus and Calender have in mind. The problem of course is precisely the fact that one must provide a *reconstruction* of Perrin's reasoning, and while Salmon and Achinstein both agree that it constitutes an argument for scientific realism, they of course disagree as to exactly what sort of argument it is—and Psillos (2011c) of course disagrees with both Salmon and Achinstein. Moreover, some philosophers, for example van Fraassen (2009), have even proposed anti-realist reconstructions of Perrin's reasoning, taking him to be only arguing for the empirical adequacy of the atomic hypothesis. The point then is that such reconstructions rarely wear their philosophical credentials on their sleeves, which leaves Salmon, Achinstein, Psillos, and indeed van Fraassen, in an awkward situation: one cannot simply present a piece of first-order scientific reasoning and declare it to *ipso facto* constitute an argument for one's preferred philosophical disposition without straightforwardly begging the question; but once one attempts to justify one's interpretation, by appealing to features of the general argumentative scheme of which it is an instance, one is back into precisely the kind of wholesale reasoning that one was trying to avoid.

The challenge then is to find some way of articulating an appropriately localised, retail argument for (or against) scientific realism that manages to provide a genuinely philosophical dimension to the first-order scientific reasoning with which we began, but without thereby stumbling into the sort of generalised considerations to which the base-rate fallacy applies. In what follows I will examine three such proposals in the contemporary literature, which while clearly far from exhaustive, do seem to cover a goodly portion of the logical space.

The first is due to Stathis Psillos, who rather than imposing a philosophical interpretation upon our first-order scientific reasoning, argues that it can nevertheless directly impact upon our second-order philosophical reflections in a way that motivates and legitimises a specific and context-sensitive approach to the scientific realism debate. The basic problem with Psillos' proposal however—or so I will argue—is that on his understanding of the dimensions of the debate, there simply is no second-order philosophical evidence for our first-order reasoning to influence.

The second proposal is due to Juha Saatsi, who rather than imposing a philosophical interpretation upon our first-order scientific reasoning as a whole, attempts to identify some distinctively philosophical elements within that enterprise. This offers an improvement over Psillos' account, which struggles to maintain the distinction between first-order and second-order evidence given the constraints imposed by the threat of the base-rate fallacy, but nevertheless similarly struggles to maintain its own distinction between our scientific reasoning and our philosophical reasoning. In both cases then, Psillos and Saatsi fail to identify any philosophical considerations which are not already part of the first-order scientific reasoning with which the scientific realist is concerned.

In the next section I consider a third proposal due to Kyle Stanford, who although not directly concerned with the problem of the base-rate fallacy, does defend a version of anti-realism with a more theory-specific orientation. More specifically, whereas Psillos and Saatsi attempt to inflate our first-order scientific reasoning into something with philosophical clout, Stanford attempts to narrow down our second-order philosophical reflection into something that makes closer contact with the specifics of our scientific theories. The general problem however remains the same—for while Psillos and Saatsi fail to make our first-order scientific reasoning general enough to make a distinctively philosophical contribution, Stanford fails to make our second-order philosophical reasoning particular enough to avoid concerns over our underlying base-rates. Either way then, I conclude, there remains a serious tension in the idea of a retail argument regarding scientific realism as presently articulated.

3.1. Psillos on first-order and second-order evidence

As Psillos (2009, pp. 75–77; see also his 2011a) now sees it, the point of the scientific realism debate is not so much to legislate over the appropriateness of our scientific methodology, but rather to attempt to *balance* the often-competing evidence that we may have regarding the approximate truth of our scientific theories. So on the one hand, there will be the first-order evidence that the scientists themselves principally take into account: the extent to which the theory is confirmed for instance, as well as other theoretical virtues such as simplicity, scope and coherence with the rest of our existing scientific worldview. While on the other hand, there will be the second-order evidence with which the philosopher is principally concerned: not so much with the specifics of the scientific theory as with the general reliability of the scientific method, its historical track-record, and the threat of underdetermination. In the case of Perrin then for example, we have various first-order considerations for accepting the atomic theory of matter—namely the surprising agreement amongst our different methods for calculating the value of Avogadro's number—while at the same time we face the competing second-order evidence—as stressed by phenomenologist dissenters to the atomic hypothesis at the time—that even our most successful scientific theories are eventually abandoned as false (cf. Psillos, 2011c).

The principal question for the scientific realist then is whether or not we can marshall any first-order evidence for a scientific theory that can also *override* the second-order evidence we may have against it. Consider for example Psillos' (1999, pp. 101–114) general response to the pessimistic meta-induction. The thought here is that while the history of science and the continual overhaul of our theoretical framework provides second-order evidence for doubting our current scientific theories, the sort of first-order evidence that we can present in their favour is in actual fact rather specific to the theories in question—it is these *particular* theoretical posits that have explanatory strength, these *particular* structural relationships that are predictively successful. Consequently, since our first-order evidence for a scientific theory is essentially unique, there can be no inductive basis for a general pessimism over their historical stability, and our second-order evidence undermined.

The idea then is that a retail argument for scientific realism need not consist of simply presenting our first-order scientific reasoning and declaring it to simultaneously specify its second-order philosophical assessment; rather, we are to maintain the distinction between the localised considerations presented by scientists, and the abstract generalisations offered by philosophers, but to reconsider the relationship between the two—specifically, that while general reflection over the reliability and historical track-record of our scientific reasoning may influence our assessment of our first-order evidence, so too can our first-order evidence influence what sort of general philosophical reflections are deemed legitimate.

The distinction however is illusory. Recall that on Psillos' account, we are to sharply individuate both the content of our scientific theories *and* the inferential methods by which we arrived at such theories. The predictive success of a scientific theory is to be attributed to specific theoretical posits and structural relationships such that there can be no general inductive assessment of their historical stability; and the reasoning by which we arrived at these theories is to be taken as thoroughly context-specific so as to avoid any worries relating to the base-rate reliability of these methods. But if that is the case—if both our methods of reasoning and the content of our scientific theories is to taken as thoroughly particular—then there will be absolutely nothing in general that the scientific realist can say about the approximate truth of our scientific theories, and thus nothing in the way of second-order evidence with which our first-order evidence can interact. Rather than articulating a way in which our first-order scientific evidence can directly influence our second-order philosophical reflections, Psillos in fact rules out the possibility of there being any second-order philosophical reflection at all; his strategy is then no different from that with which we began, and threatens to reduce the scientific realism debate to nothing more than mere repetition of the first-order scientific reasoning with which it was supposed concerned.

3.2. Saatsi on material postulates

A more promising proposal is due to Saatsi (2010), who distinguishes between what he calls *form-driven* arguments for scientific realism and *content-driven* arguments for scientific realism. The distinction is more-or-less equivalent to that between Magnus and Callender's distinction between wholesale and retail arguments—the traditional no-miracles argument is both a wholesale argument in that it concerns our scientific methodology in general, and a form-driven argument in that it is concerned with the reliability of inference to the best explanation in the abstract—but is logically distinct in that it concerns structural features of the argument in question rather than just its intended scope. More specifically, a content-driven argument for scientific realism will be concerned with what Norton (2003) calls the material postulates of that argument i.e. the assumptions of uniformity upon which the argument depends.

The idea is best illustrated in the case of induction, where rather than attempting to resolve the formidable questions as to whether or not extrapolation to the future is a reliable method of inference, we can at least make some progress attending to the somewhat more tractable issue as to what sorts of extrapolations are going to be more reliable than others. The problem of induction remains unresolved, but we can still ask whether or not *these* samples constitute a good inductive basis. Are they representative? Do they really support regularities? Similarly then in the case of scientific realism: the philosopher of science should not be concerned with whether or not inference to the best explanation is a generally reliable method of inference, or whether or not it can be given a non-circular justification to the committed sceptic, but whether or not these particular features of a theory are of the right kind upon which any reasonable inference to the best explanation can be based.

We can think of Saatsi's proposal then as an attempt to find a third way between the two extremes of offering a second-order philosophical argument for scientific realism that falls foul of the base-rate fallacy, and offering a first-order scientific argument for scientific realism that simply begs the question over its interpretation. The idea is that the distinction between first-order evidence and second-order evidence, and between our scientific reasoning and our philosophical reasoning, does not line up in the way assumed throughout this paper—and that in fact there exists some genuinely philosophical aspects to the first-order evidence we may have for the approximate truth of our scientific theories, in this case concerning the material postulates upon which that reasoning depends. As Saatsi sees it, while an approximately true scientific theory must *ipso facto* have latched onto the right sorts of material postulates, this will not always be made explicit by the scientific methods used to arrived at that theory; hence

> If a scientist appeals to a theory T because it is the simplest and the most unifying, and hence the most explanatory perhaps, it is a task for the philosopher to make explicit how these contextual judgements reflect the particular material facts, given the scientific background knowledge of the domain in question. Only once material postulates have been made transparent can we compare them with the particular assumptions underwriting some commensurate inductions to the observable. (ibid.: 26)

The picture then seems to be something like the following. Scientists give their usual first-order reasoning for accepting a particular scientific theory. The philosopher of science then considers the material postulates underlying this reasoning, essentially those features of the world that the scientist is implicitly assuming to be uniform when making his inference. This does not constitute a second-order reflection upon the reliability of our scientific practices as a whole—and so is not at the mercy of unknowable base-rates—but is rather concerned with the specifics of the scientific theory in question. The question however is the extent to which this assessment of material postulates offers a genuine philosophical contribution to our assessment of that theory.

The problem is that the assessment of our material postulates appears to be part and parcel of our first-order scientific reasoning. When scientists come to accept a particular scientific theory, one of the things they take into account will be the sorts of assumptions of uniformity upon which the theory depends. Indeed, most of these assumptions will themselves be the results of previously accepted scientific theories—Perrin's work on atoms for example crucially depended upon the assumption of Brownian motion in gases, a material postulate for which, in turn, various first-order scientific evidence had been offered in support. So if Saatsi is right, and if the scientific realism debate really does constitute a distinctive philosophical contribution to the question over the approximate truth of our scientific theories, then there must be some considerations *over and above* our first-order scientific practices for accepting these material postulates to be sound.

There seems to me to be two ways one can go here.[4] One option would be to offer arguments for the claim that—in general—the inferences that we make on the basis of this or that material postulate are reliable; the problem of course is that then we would just be making a second-order assessment of our first-order practices, and must again attend to the problem of unknown base-rates. Just because our current scientific theories are successful and depend upon a particular class of material postulates, it does not follow that any arbitrary inference based upon that class of material postulates will be reliable: just as with the traditional argument from predictive success to truth, there may be an enormous number of *unsuccessful* inferences based upon those material postulates that our limited sampling overlooks. The other option would be to focus the philosophical dispute onto the question as to *which* material postulates

[4] I would like to thank an anonymous referee for pushing me repeatedly on this point, and for helping me to better understand what a first-order philosophical argument for scientific realism could look like. I regret that I am still unable to answer their concerns in a fully satisfactory manner, but appreciate their help in sharpening my thinking on this point.

are being assumed. In the case of Perrin again for example, the scientific realist may argue that the material postulates at work are various principles of uniformity holding between the macroscopic and the microscopic, such that the behaviour of atoms can be understood on a par with the behaviour of larger objects that accelerate and decelerate through collisions and the conservation of momentum; the empiricist by contrast may argue that the principle is more attenuated—that the microscopic world behaves *as if* it bumped and banged its way around like the observable world, or some story to that effect. The philosophical debate would then be over the relative merits of these competing accounts. But now we seem to be back where we started, having exchanged the *assessment* of our first-order scientific practices for the *interpretation* of our first-order scientific practices.

There is undoubtedly more to be said here, but the general dilemma seems to remain—any argument for scientific realism that attends to the content of the inference in question will either be so theory-specific as to collapse into our first-order scientific reasoning, or operate at a level of epistemological generality that threatens to commit the same base-rate fallacy all over again.

4. The pessimistic meta-induction and Stanford's unconceived alternatives

The standard argument offered against scientific realism is a pessimistic meta-induction based upon the history of our scientific practices. The basic idea is that since our most successful scientific theories in the past have all turned out to be false, we have enumeratively inductive reasons for supposing that our currently successful scientific will also eventually be abandoned; the basic problem with such reasoning is that, just as with the scientific realist's inference from success to truth, the strength of the argument depends upon facts about the base-rates of our overall sample space that the anti-realist is not in a position to know independently of having already endorsed his conclusion.

There are two ways to illustrate this worry, one concerning the overall distribution of failed theories across the history of science; and a more general concern as to the overall ratio of failed scientific theories at any particular moment of time. In the first case we can simply note that it is perfectly consistent for the history of science to furnish us with an overwhelmingly large number of unsuccessful theories, yet for the vast majority of these failures to be confined to one narrowly defined domain of inquiry, or to one narrowly defined period of time. In such a situation, we would hardly be justified in inferring the probable falsity of an arbitrary scientific theory, let alone the probable falsity of our contemporary scientific theories in general, just because (say) eighteenth-century biology was extremely unstable. This would be an example of what Lange (2002) identifies as a *turnover fallacy*, noting that what the pessimistic meta-induction requires is the temporally specific evidence to the fact that at most past moments of time, most of the theories at that time were false (ibid.: 284).

The more general problem can be illustrated by simply supposing that when we come to assess our scientific theories, the number of false scientific theories massively outnumbers the number of true scientific theories. In such a situation, the number of predictively successful scientific theories that nevertheless turn out to be false might be extremely high, even though the *likelihood* of any particular predictively successful theory being false is extremely small. As Lewis (2001) argues, this would be consistent with the anti-realist's historical evidence, but would in fact undermine the anti-realist's conclusion that our currently successful scientific theories are most likely false. As with the no-miracles argument in favour of scientific realism then, the pessimistic meta-induction for scientific anti-realism also depends fundamentally upon the base-rate probability of an arbitrary scientific theory being true, which is of course the precisely the issue at stake.[5]

Thus just as advocates of scientific realism have subsequently turned their attention to more specific arguments for their position, so too have critics of scientific realism reformulated their approach. Although he does not address the issue of the base-rate fallacy explicitly, Stanford (2006) has recently defended an instrumentalist position in the scientific realism debate which, rather than restricting our beliefs on the basis of a general distinction between e.g. the observable and unobservable content of our scientific theories, proposes instead that those parts of our scientific theories that are to be believed will vary from context to context; the idea then is that the difference between the realist and the instrumentalist will be "a local difference in the specific theories each is willing to believe on the strength of the total evidence available" (ibid.: 205).

Stanford's argument is a variant on traditional anti-realist arguments, seeking to combine the strengths of both the pessimistic meta-induction and the problem of underdetermination. He argues that what the history of science shows us is not that our past scientific theories have all turned out to be false, but rather that our past scientific theories have all been underdetermined by the evidence at that time—even if their theoretical alternatives were unconceived by the scientists themselves and only recognisable in retrospect. His resulting instrumentalism is then based upon which specific claims of our scientific theories seem to be more or less vulnerable to this risk of underdetermination: claims about the behaviour of everyday medium sized objects—what other anti-realists would designate as the 'observable' content of our scientific theories—have enjoyed enough epistemological stability throughout the history of science to warrant our belief; but crucially, so too have various well-entrenched 'theoretical' or 'unobservable' claims about certain aspects of the world. The precise point at which we draw the line will therefore vary from theory to theory, since just because the historical record gives us reasons to suppose that a great deal of the claims of one particular scientific theory are likely to be underdetermined by the evidence, it does not follow that the historical record will give us reasons to be similarly sceptical of another scientific theory.

What is particularly interesting for our purposes is how Stanford's position offers a way for the anti-realist to avoid the problems of base-rates associated with traditional arguments against scientific realism. Unlike the familiar problem of underdetermination, Stanford's argument does not trade on the mere logical possibility of constructing empirically equivalent alternatives to our scientific theories, but rather seeks to demonstrate through concrete historical investigation that, at any particular moment of time, there were unconceived alternatives to our accepted scientific theories that were in fact eventually adopted by the scientific community (ibid.: 20–21). In the case of the traditional pessimistic meta-induction, the argument was based upon the eventual failure of previously accepted scientific theories, and was therefore vulnerable to the retort that our contemporary theories are sufficiently different—more mature, greater predictive power, etc.—from their predecessors as to undermine any inductive generalisation; or in terms of our current framework, that the *probability* of an arbitrary scientific theory being false is not a good guide to the *likelihood* of our current scientific theories being false.

[5] As Saatsi (2005, p. 1097) points out of course, there is a weaker reading of the pessimistic meta-induction whereby one simply presents instances of historical failure as a series of counterexamples to the scientific realist's inference from success to truth, without attempting to draw any general conclusions on their basis. Such reasoning however clearly falls far short of providing a positive argument for anti-realism.

Stanford's argument neatly side-steps these issues since it does not concern our overall sample of scientific theories, and thus does not depend upon the base-rate probability of an arbitrary scientific theory being true; rather, Stanford's argument targets *the scientists themselves*, arguing that there is no corresponding reason to suppose contemporary scientists to be any better at exhausting the space of relevant alternatives than their forbearers, and therefore that a degree of instrumentalist caution is advised (ibid.: 44–45).

Unlike the proposals due to Psillos and Saatsi therefore, who in effect attempt to show that our first-order scientific reasoning can itself make a philosophical contribution without straying into base-rate territory, Stanford's position attempts to deflate our second-order philosophical reasoning into something particular enough to side-step concerns regarding base-rates. There are however a number of concerns that can be raised regarding Stanford's position. The first is that it is far from clear why we should suppose that contemporary scientists are no better off at exhausting the space of relevant alternatives than earlier scientists. Indeed, if we are going to concede that our contemporary scientific theories are in some way superior to their earlier formulations sufficient to problematise a traditional pessimistic meta-induction, it would seem odd not to suppose that at least some of this theoretical progress can be attributed to cognitive advances in the methods of the scientists themselves (Psillos, 2009, pp. 72–75).

The deeper problem with Stanford's account however again turns upon the intended scope of his position, and there is in fact something of a tension here. On the one hand, the argument is definitely supposed to move away from the sweeping generalisations that have characterised the scientific realism debate, and to target instead the specifics of individual scientific theories: not only does Stanford begin with a detailed case study of three important episodes in the history of biology, but the instrumentalist conclusion that he draws maintains that the appropriate epistemic attitude that we hold towards a scientific theory will depend upon the specific history of that domain and will vary from context to context. Yet on the other hand, the engine driving Stanford's argument is still a perfectly general inductive inference—on the basis of the fact that previous scientists failed to exhaust the space of alternative theories in the past, we have reasons to suppose that contemporary scientists have similarly failed to exhaust the space of alternative theories in the present and that our current theories are also radically underdetermined. Yet while the threat of unconceived alternatives may well be more robust than the threat of the standard pessimistic meta-induction, and while it may indeed avoid the need for knowing the underlying base-rate of any arbitrary scientific theory being true, it seems that for the inference to be compelling we must simply presuppose a different base-rate concerning the reliability of scientists: maybe our historically unconceived alternatives are all heavily biased towards researchers working in a very specific domain of inquiry; or maybe while the history of science furnishes us with a great number of instances of scientists failing to exhaust the relevant possibilities, this is due to the relatively large number of scientific practitioners, rather than the likelihood of any arbitrarily successful researcher failing to consider another alternative theoretical formulations.

So the same general dilemma recurs. If Stanford wants his historical studies to show us anything in general about the philosophy of science, he must operate at a level of generality in which the issue of base-rate probabilities re-emerge, even if it concerns the underlying reliability of scientists rather than theories; and if he remains at the level of individual scientific theories and individual scientists, then he will be merely engaging in more of the same first-order reasoning about the evidence for our theories in which scientists already engage.

5. Naturalism and normativity

In the early half of the twentieth century, the scientific realism debate was primarily concerned with the logico-semantic structure of a scientific theory: with whether or not our putatively 'theoretical' discourse was to be taken at face-value and thereby entailing ontological commitment to various unobservable entities and processes, or if it was to be somehow reinterpreted or eliminated as a purely syntactic device. On such a construal, the intended scope of the scientific realism debate was extremely clear-cut, since one could *appeal* to our actual scientific practice in adjudicating between these competing semantic claims—the open-endedness of scientific research, for example, weighed heavily against any attempt to explicitly define the theoretical vocabulary in terms of our observational vocabulary—but without thereby *engaging* in that practice. But with the demise of logical empiricism, the contours of the scientific realism debate shifted to a more epistemological orientation, and here the crucial distinction between our first-order scientific reasoning, and our second-order philosophical reasoning about our first-order scientific reasoning, began to blur. The idea that the philosopher of science somehow has access to a deeper or more profound source of knowledge, on the basis of which he can sit in judgement over our first-order scientific practices, is one that has been firmly rejected by contemporary analytic philosophy; indeed, if there has been one guiding intuition framing the contemporary scientific realism debate above all others, it is that any such investigation must proceed in conjunction with a suitably naturalistic methodology—that is, in acknowledgement of the fact that our philosophical investigations are *continuous* with our empirical investigations, and do not constitute some higher court of epistemic appeal.

Yet arguably, it is precisely this endorsement of a naturalistic methodology that has led the scientific realism debate into its current impasse. If the philosophy of science does not constitute an independent source of reasons and arguments concerning the approximate truth of our scientific theories as the naturalist contends, then our second-order reflections can differ only in degree from our first-order scientific reasoning; but the most obvious way in which to understand this difference, between the *specific* first-order evidence assembled by the practising scientist as opposed to the *general* second-order evidence assembled by the philosopher, has been explicitly undermined by worries relating to the base-rate fallacy and an increasing tendency to direct one's philosophical attention to extremely context-specific instances of our scientific methodology. The problem seems to be then in maintaining *both* that the philosophy of science should be concerned with individual scientific theories rather than with general methodological pronouncements, *and* that the philosophy of science should be understood as continuous with our scientific methods and not as contributing a distinctive source of normative evaluation. The two taken together threaten to squeeze the scientific realism debate into providing nothing more than the superfluous repetition of our first-order scientific practices.

For some at least, this is good news, and the scientific realism stands exposed (again) as the philosophical folly that it is. What is interesting to note however is that unlike other well-known pronouncements of the death of scientific realism, this particular result appears to be entirely self-generated. Carnap (1950, 1974) of course famously took the debate over the existence of electrons and neutrinos to be a philosophical pseudo-problem: one simply took a pragmatic decision as to whether or not one wished to use a language that included terms like 'electron' and 'neutrino', and declared this to be as close as one could get to explicating an otherwise intractable metaphysical muddle. More recently, van Fraassen (2002) has argued that one should not understand realism and empiricism as substantive philosophical positions backed up with considered arguments, but rather as the expression of different epistemic standards ultimately based upon our cognitive values and aesthetic sensibilities. But both positions are philosophical controversial, and which one may come to hold or reject quite independently of one's views in the philosophy of science: to endorse Carnap's meta-ontological dissolution of the scientific realism debate, one must also endorse a series of further claims regarding for example the distinction between the analytic and the synthetic; and to follow van Fraassen, one must subscribe to an extremely permissive conception of rationality whereby anything that is not strictly forbidden is permissible, and where there is no universal standard for adjudicating between internally consistent epistemic stances. The conclusion argued for in this paper however is different: if the philosophy of science is understood to be continuous with our scientific practices, and if the philosophy of science cannot operate on a more general level of evaluation without risk of statistical error, then there just doesn't seem to be anything left for the scientific realism debate to be about.[6]

One possibility left open by the preceding discussion would be to recast the scientific realism debate as concerned with something like the deep metaphysical structure of reality, a level of description that is understood to lie outside of the purvey of our empirical investigations. The idea then would be that while our first-order scientific methods are indeed highly reliable in determining the regularities that hold among the surface phenomena, there still remains the further question as to whether or not our scientific descriptions have really latched onto the fundamental structure of the world—whether they have carved nature at its joints, as it were. The role for the philosophy of science would then be to provide just this level of reassurance—to confirm for instance that electrons and neutrinos really are the sort of natural kinds with which our first-order scientific practices should be concerned. This would certainly legitimise a distinct role for the philosophy of science independently of our first-order scientific practices; and it would be the sort of philosophical contribution that could be made on a case-by-case basis, without sweeping generalisation or risk of statistical error.

It is also—I take it—an extremely unattractive option for the contemporary scientific realist, and not something endorsed by Psillos, Saatsi or Magnus and Callender. It would of course constitute an explicit rejection of the naturalistic methodology upon which most of the contemporary scientific realism debate is based, as well as presupposing an unrealistic conception of the scope and authority of the philosophy of the science. The only other option then would be to rethink what exactly the normative dimension of the philosophy of science amounts to. Questions as to whether or not we should accept a particular scientific theory seem best left to the scientists themselves; but this is not the only normative issue that arises within our scientific practices, and nor should it be the only one of interest to traditional scientific realists.

6. Conclusion

Recent literature in the scientific realism debate has been rightly concerned with the extent to which the principle arguments that frame the issue—most prominently, the no-miracles argument and the pessimistic meta-induction—are vulnerable to a particular species of statistical fallacy. Responses to this worry have therefore pursued an increased emphasis upon the specifics of individual scientific theories, the approximate truth of which will not depend upon the reliability of our scientific methodology considered in general. However, this emphasis upon the first-order scientific evidence that we may have for a scientific theory, in conjunction with the wide-spread naturalistic conviction that the philosophy of science does not provide an independent source of evaluation, threatens to eliminate the possibility of any distinctively philosophical contribution to his debate. The real challenge of the base-rate fallacy therefore is not so much a change of emphasis for the philosopher of science, but a serious reconceptualisation of what the normative dimensions of the scientific realism debate can amount to.

Acknowledgements

Earlier versions of this paper were given at the Society for Exact Philosophy in Montréal, and at the British Society for the Philosophy of Science in Exeter. I would like to thank the organisers and participants of these conferences, and in particular Juha Saatsi and Ioannis Votsis for their comments and feedback. Thanks also to Stathis Psillos for looking at a rough draft of this paper, Michaelis Michael for extended conversations, and an extremely helpful anonymous referee for this journal for sharpening up many of my arguments. Finally, I would like to thank Hannes Leitgeb and the Munich Center for Mathematical Philosophy, and the Alexander von Humboldt Foundation, for supporting this research.

References

Achinstein, P. (2001). *The book of evidence*. New York: Oxford University Press.
Boyd, R. (1980). Scientific realism and naturalised epistemology. *Proceedings of the Biennial Meeting of the Philosophy of Science Association, 2*, 613–662.
Carnap, R. (1950). Empiricism, semantics and ontology. *Revue Internationale de Philosophie, 4*, 20–40.
Carnap, R. (1974). *An introduction to the philosophy of science*. New York: Basic Books.
Fine, A. (1984). The natural ontological attitude. In J. Leplin (Ed.), *Scientific realism* (pp. 83–107). Berkeley: University of California Press.
Fine, A. (1986). Unnatural attitudes: Realist and antirealist attachments to science. *Mind, 95*, 149–177.
Howson, C. (2000). *Hume's problem*. New York: Oxford University Press.
Lange, M. (2002). Baseball, pessimistic inductions and the turnover fallacy. *Analysis, 62*, 281–285.
Lewis, P. (2001). Why the pessimistic induction is a fallacy. *Synthese, 129*, 371–380.
Lipton, P. (2000). Tracking track-records. *Proceedings of the Aristotelian Society, 74*, 179–205.
Magnus, P. D., & Calendar, C. (2004). Realist ennui and the base rate fallacy. *Philosophy of Science, 71*, 320–338.
Musgrave, A. (1988). The ultimate argument for scientific realism. In R. Nola (Ed.), *Relativism and realism in science* (pp. 229–252). Dordrecht: Kluwer.
Norton, J. D. (2003). A material theory of induction. *Philosophy of Science, 70*, 647–670.
Psillos, S. (1999). *Scientific realism: How science tracks truth*. London: Routledge.
Psillos, S. (2009). *Knowing the structure of nature: Essays on realism and empiricism*. Basingstoke: Palgrave Macmillan.
Psillos, S. (2011a). Moving molecules above the scientific horizon: On Perrin's case for realism. *Journal for General Philosophy of Science, 42*, 339–363.
Psillos, S. (2011b). Choosing the realist framework. *Synthese, 190*, 301–316.
Psillos, S. (2011c). Making contact with molecules: On Perrin and Achinstein. In G. Morgan (Ed.), *Philosophy of science matters* (pp. 171–191). New York: Oxford University press.
Putnam, H. (1975). What is mathematical truth? In *Mathematics, matter and method, philosophical papers* (Vol. I, pp. 60–78). Cambridge: Cambridge University Press.
Saatsi, J. (2005). On the pessimistic induction and two fallacies. *Philosophy of Science, 72*, 1088–1098.
Saatsi, J. (2010). Form vs. content-driven arguments for realism. In P. D. Magnus & J. Busch (Eds.), *New waves in philosophy of science* (pp. 8–28). Basingstoke: Palgrave Macmillan.
Salmon, W. (1984). *Scientific explanation and the causal structure of the world*. Princeton: Princeton University Press.
Stanford, P. K. (2006). *Exceeding our grasp: Science, history and the problem of unconceived alternatives*. New York: Oxford University Press.
van Fraassen, B. C. (1980). *The scientific image*. Oxford: Clarendon Press.
van Fraassen, B. C. (1989). *Laws and symmetry*. Oxford: Clarendon Press.
van Fraassen, B. C. (2002). *The empirical stance*. New Haven: Yale University Press.
van Fraassen, B. C. (2009). The perils of Perrin, in the hands of the philosophers. *Philosophical Studies, 143*, 5–24.

[6] This conclusion may in fact be similar to Fine's (1984, 1986) assessment of the scientific realism debate, and his so-called Natural Ontological Attitude. For on the one hand, Fine complains that both realists and empiricists go wrong in trying to add unnecessary philosophical complications to our otherwise perfectly well-understood scientific practices. However on the other hand, Fine seems to understand these complications as predominantly semantic and concerning competing accounts of truth, as opposed to the broadly methodological considerations presented here. I am therefore unsure as to the exact parallels to be drawn here.