Chapter 4
Miracles and Structural Realism

John Worrall

4.1 Introduction

The often breathtaking predictive success of *some* theories in contemporary science inclines most of us towards scientific realism: surely what those theories say about the 'unseen world' lying 'beyond the phenomena' must be at least approximately correct if they can score such dramatic, empirically checkable, successes? The facts about theory-change in science, on the other hand, seem to speak in favour of an anti-realist view: scientists have in the past held theories that were also dramatically predictively successful and yet which are now 'known to be false' (because they are inconsistent with our latest theories). Given this, what guarantee can there possibly be that our latest theories will not themselves be rejected and replaced by quite different ones at some time in the future? And if so, how can we reasonably hold that our current theories are true? And when we think about how radical some of those theory-changes appear to have been, how can we hold that our current theories are likely even to be *approximately* true? My [21] argued that, although these two much-discussed considerations thus seem to pull sharply in opposite directions, they can in fact be reconciled within a version of realism—namely, *structural realism*.

The first, apparently pro-realist, consideration has often been developed as 'the no-miracles argument' (hereafter the NMA, or rather, as we shall see, 'the' NMA). The intuition is roughly that it would be a miracle if current scientific theories enjoyed the predictive successes that they do if what they claimed was going on 'behind' the phenomena were not at least approximately correct; but we should not accept that miracles have occurred unless there is no non-miraculous alternative; and here the (approximate) truth of what the theories say about the 'noumenal' world is exactly a non-miraculous alternative explanation of their empirical success. The second, apparently anti-realist, consideration has often been developed as 'the pessimistic (meta-) induction' (hereafter, the PI). Roughly: theories that were accepted in the past (exactly on the basis of the predictive success emphasised by the

J. Worrall (✉)
Philosophy, Logic & Scientific Method, London School of Economics, London, UK
e-mail: j.worrall@lse.ac.uk

E.M. Landry, D.P. Rickles (eds.), *Structural Realism*, The Western Ontario Series in Philosophy of Science 77, DOI 10.1007/978-94-007-2579-9_4,

NMA) have subsequently turned out to be (perhaps radically) false; so, we should infer (inductively/probabilistically) that our current theories are (perhaps radically) false too.

However, Colin Howson has argued (see [6], chapter 3) that the NMA—in so far as it can be turned into a precise argument at all—in fact embodies an elementary probabilistic fallacy: often called 'the base rate fallacy'. While, ironically enough, Peter Lewis [11] has (independently) argued that essentially the same fallacy underlies, and therefore vitiates, its seeming-competitor argument—the PI.

If these arguments really do establish that the NMA and the PI are fallacious, then this would seem to destroy the basic problematic at which structural realism is addressed. And Magnus and Callender have indeed recently urged that, since 'the major considerations for and against realism come to naught', the whole scientific realism debate (at any rate in what they call the 'wholesale' sense) should be 'dissolved', as resulting in nothing but 'ennui', [14, pp. 321–322].

Is there anything in these recent arguments that should concern the structural realist or force her into a state of terminal ennui? In this paper I restrict myself to the concerns about 'the' NMA (leaving those about 'the' PI for another occasion [26]). I argue that the 'difficulties' raised in the literature are no more than artefacts of the (misguided) way in which the considerations underlying 'the' no miracles argument have been formalised, or 'modelled', as precise arguments. The underlying intuition remains untouched and remains a good (though, of course, far from conclusive) reason for adopting structural realism.[1]

4.2 No Miracles Reconsidered: The Intuitions

Consider a classic, and by now well-worn, example that elicits the intuitive 'no miracles response' (at any rate in yours truly). Fresnel's theory states that light consists of (not directly observable) waves transmitted through a (not directly observable) all-pervading elastic medium. His theory turned out to entail, as Poisson demonstrated but as Fresnel himself had never suspected, the directly empirically checkable result that if a small opaque disc is held in light diverging from a point source and if the 'geometric shadow' of the disc (that is, the area of complete darkness that would exist if the laws of geometric optics were strictly correct) is carefully examined, then the centre of that 'shadow' will in fact be seen to be illuminated, and indeed just as strongly illuminated as if no opaque disc were present. Although most of the French Academicians thought that this was a clear-cut *reductio* of the theory, when

[1] There have also been a number of direct criticisms of Structural Realism in the recent literature, many of them based on what might be called the 'Newman objection' (see [15], the revival of that argument in [2], and more recently [9]). These criticisms are also not dealt with in the present paper but *are* addressed and rebutted in [25].

Arago performed the experiment it turned out that the 'white spot' does indeed, and contrary to all prior expectations, exist.[2]

Whatever esoteric philosophical considerations may be raised, it is difficult to resist the feeling that if a theory can make such a striking, seemingly improbable prediction that nonetheless turns out to be empirically correct, then the theory must somehow be 'approximately true'—it must have somehow latched on, no doubt in an approximate (but nonetheless substantial) way, to the 'deep structure' of the universe: to how things really are in the 'noumenal world' behind or beyond the phenomena. Duhem, who was not the instrumentalist he is often considered but rather a structural realist, put it eloquently [3, p. 28]:

> The highest test, therefore, of our holding a classification as a natural one is to ask it to indicate in advance things which the future alone will reveal. And when the experiment is made and confirms the predictions obtained from our theory, we feel strengthened in our conviction that the relations established by our reason among abstract notions truly correspond to relations among things.

A theory gives us a 'natural classification', according to Duhem, just in case 'the relations' it posits 'truly correspond to relations among things'. Our 'conviction' that Fresnel's theory represents such a natural classification is 'strengthened' because it would, it seems, be extremely unlikely that the theory could have such a striking, and empirically correct consequence (that it should 'indicate in advance things which the future alone [revealed]'), were what it said about the reality that underlies phenomena such as the 'white spot' *not* a natural classification, that is, not in some sort of (approximate) correspondence with reality.

Or take another much-discussed example: Quantum Electrodynamics predicts the magnetic moment of the electron to better than one part in a billion. (Theory yields $1.00115965246 \pm 0.0000000002$ as the value, while, early in this century, the most sensitive observations yielded the value of $1.00115965221 \pm 0.00000000003$!).[3] Again it seems difficult to resist the feeling that the theory *must* somehow have latched on to the way things are 'underneath' the phenomena if it can get such a prediction correct to such an implausibly high degree of accuracy.

The talk about theories being 'approximately true' or (better) 'somehow latching on to the way that things are beneath the phenomena' may be imprecise but it is clearly necessary. This is emphasised by the considerations underlying the PI—later (better) theories tell us that Fresnel's theory is strictly false—but the point is independent of that argument. No one believes, even ahead of any further 'scientific revolution', that Quantum Field Theory is true—indeed there are questions about whether a fully coherent theory can be articulated at the present time. No one even believes (or ought to believe) that Quantum Mechanics itself, for all its stunning success, will survive entirely unscathed. (Its two basic postulates are clearly mutu-

[2] The real history, as I show in my [20], was a good deal more interesting and a great deal less clear-cut. However the real historical details, although they do centrally affect the issue of what counts as a successful prediction (and why predictions carry more confirmatory weight), do *not* affect the philosophical issue about the link between successful prediction and realism.

[3] Values quoted from [6].

ally incoherent; and it also fails to cohere with the General Theory of Relativity.) So the realist claim must in general be that it would be a 'miracle', not if the theory at issue failed to be 'outright' true, but rather if it failed to be somehow *approximately* true.

Whether the NMA can be given some more exact construal will be our central topic. But it does seem clear that science centrally embodies the underlying intuition and would not be possible if it did not. This is reflected in the fact that appeal to the intuition is implicit, not just in the justification of observation-transcendent theories, but also in the justification of the standard empirical generalisations that everyone—including those who are anti-realist about scientific theories—accepts. This has gone largely unacknowledged in the recent literature, but it *was* recognised, and emphasised, by Poincaré, who wrote [16, pp. 149–150]:

> We have verified a simple law in a considerable number of particular cases. We refuse to admit that this coincidence, so often repeated, is a result of mere chance and we conclude that the law must be true in the general case.
>
> Kepler remarks that the positions of the planets observed by Tycho are all on the same ellipse. Not for one moment does he think that, by a singular freak of chance, Tycho had never looked at the heavens except at the very moment when the path of the planet happened to cut that ellipse … [I]f a simple law has been observed in several particular cases, we may legitimately suppose that it will be true in analogous cases. To refuse to admit this would be to attribute an inadmissible role to chance.

It would, Poincaré is saying, constitute an incredible coincidence—a 'miracle', if you like—if Kepler's simple (first) law were instantiated by all the planetary positions that had so far been checked but was not in general true (that is, not instantiated also by all the—past and future—unobserved planetary positions). No instrumentalist or constructive empiricist known to me fails to endorse the acceptance as rational of standard empirical generalisations (again generally, as the Kepler case illustrates, as approximately, rather than outright, true)—on the basis of what is of course bound to be a finite set of actual observations. They are therefore all relying—as Poincaré argues—on exactly the no miracles consideration that they vigorously deny should be thought of as persuasive when it comes to observation-'transcendent', theoretical claims.[4]

The intuition underlying the no miracles argument also underwrites persuasive arguments in a variety *both* of scientific *and* of more commonplace circumstances. Maxwell's work initially left open the possibility that there might be two different media filling space: the optical ether *and* the electromagnetic field. But once he had discovered that waves were transmitted through the field at the velocity of light, he immediately inferred that it would be miraculous if there were two media each of which just happened to transmit disturbances at exactly the same rate; and hence he inferred that there is only one medium—the field—and that light is in fact an electromagnetic wave. Einstein refused to admit that the parameter measuring a

[4] This is why it seems disingenuous of Magnus and Callender [14] to take it that everyone agrees that we are entitled to make the standard 'horizontal' inductive inferences and then to lay this aside in their discussions of differences of opinion concerning the NMA.

body's responsiveness to an applied force (its inertial mass) and its gravitational action (its gravitational mass) could be identical by accident. It would be a 'miracle' if these two conceptually distinct quantities just happened invariably to have the same value—so some non-miraculous account must be sought and was of course found in the form of the General Theory of Relativity.

Admittedly it is easy to produce alleged 'miraculous coincidences' pretty well at will and many people have been seduced by cooked-up 'coincidences' into accepting conclusions that are themselves quite staggeringly improbable—I think in particular of arguments concerning the so-called anthropic principle and some arguments for the existence of god. So we certainly need to take care in this area. Nonetheless it seems difficult to resist the idea that there is *something* important in the intuitions underlying the truly persuasive instances of the NMA.

Philosophers—on the whole quite rightly—are, however, suspicious of 'intuitions' and try to capture what, if anything, is valuable in them in more rigorous arguments, whose credentials can in turn be examined more sharply. So how, if at all, can the intuitions elicited by the predictive success of at least some theories be captured in some more precise argument?

4.3 How *Not* to Formalise the NMA

4.3.1 The Scope of 'the' Argument

What exact scope should we expect such a formalised argument to have? Hilary Putnam [18, p. 19], suggested that we should think of scientific realism as itself a sort of 'overarching scientific hypothesis' that (allegedly) provides the 'best explanation' of the success of 'science'. This idea was developed by Richard Boyd [1] and endorsed by Stathis Psillos [17], who calls it the 'explanationist defence' of scientific realism. So the idea is that the argument is to be regarded as a grand, meta-level 'abduction' or 'inference to the best explanation'. The best (perhaps only) explanation of the success of science in general is the truth (or approximate truth) of its theories. So we are entitled to infer that its theories are indeed (at least approximately) true—that is, we are entitled to infer the thesis of scientific realism. And this inference is a scientific inference, no different in form from the particular inferences to the best explanation routinely drawn within science.

But surely a number of fundamental objections to this form of the argument were (or ought to have been) apparent from the beginning. The underlying idea is that there is some general scientific method for producing theories—one that has been so successful that we are entitled to infer at least to the approximate truth of its products. But, allowing for the moment that there may be such a method, it is certainly not uniformly successful. There is a wide variety of sciences, not all of them 'surprisingly successful', certainly not in any sense that elicits the 'no miracles intuition'. Nothing known to me in the sciences of sociology, parts of psychology, dietetics etc provides any reason to make one think that their theories have successfully penetrated to the noumenal world ‘beyond’ the phenomena. Nor does the fact that, for example, ‘scientific’ Creationism has successfully continued to attract adherents and exercise a good deal of political and social power have the slightest implication for the correctness of its theoretical claims. Nor even does the sociological fact that some theory has attracted a good deal of reasonably extended attention within ‘respectable’ scientific circles present any temptation in itself to cry ‘no miracles’. It does seem that Larry Laudan, for example, must have been using a far-too-undemanding notion of what it takes for a theory to be ‘successful’ when including such vapid ideas as the aetherial musings of Hartley and LeSage amongst those theories that ‘were once-successful but are now known to be [radically] false’ [10].

Now Putnam was of course quick to restrict the argument, from the beginning, to ‘mature’ science. But, as I have argued elsewhere [21], the only reasonable characterisation seems to be that a science achieves ‘maturity’ once its basic theories have turned out to be predictively successful. So this means in fact that we must first check whether the theories in some field have enjoyed predictive success before including them within the scope of the ‘overall’ inference. But this surely means in turn that the alleged general character of the argument evaporates on reflection: the argument is general only in so far as it consists of the union of a series of individual arguments about particular theories—arguments that infer the likely approximate truth of individual theories such as Fresnel’s or Quantum Field Theory from their particular empirical successes. Aside from this, how could any general inference to the best explanation of the success of science ever have been thought of as itself a *scientific* explanation?

There has been a good deal of talk in the explanation literature about ‘loveliness’ and the like[5]—the idea that explanatoriness is a quality that a theory may possess over and above its degree of empirical support. However, the history of science seems to me to show conclusively that which theories are found to be ‘explanatory’ is a historically contingent issue dependent on which theories have the highest degree of support. A classic instance involves action at a distance—initially branded as incomprehensible and so no possible part of any theory that could count as ‘explanatory’ (Newton himself of course shared this view), it became a perfectly acceptable, ‘explanatory’ notion because of the overwhelming empirical success of Newton’s theory: to the extent that when, for example, Coulomb came to formulate his law of electrostatic attraction and repulsion, there was no concern at all about its being an action at a distance theory. Of course later, in the light of the still better empirically supported theories of relativity which reject the idea of action at a distance, it was again abandoned as a ‘non-explanatory’ idea. What we find ‘explanatory’ or not is dependent on empirical success. And this means independent predictive success.

---

[5] See for example [13].

If scientific realism were to be considered an overarching meta-level but nonetheless scientific explanation of science's success, it would therefore, need to have *evidence* in its favour. What could that be? What, more specifically, could count as an independent predictive test of (any version of) scientific realism as a general, in Magnus and Callender's terms [14] 'wholesale', view? The only candidate would seem to be the (meta-level) 'prediction' that the next theory to be accepted in any field will be successful, that is, will enjoy independent predictive success over and above that enjoyed by the currently accepted theory. But, restricting consideration, as we already saw we must, to mature science, this will be trivially or definitionally true; and so its fulfilment cannot be regarded as an 'empirical success' on a par with those enjoyed by (ordinary 'object level') scientific theories themselves. A field becomes mature when its accepted theories are successfully predictive, and science would clearly never accept a new theory as superior to a currently accepted one unless that new theory not only scored the same successes as its predecessor, but also enjoyed predictive successes over and above those shared with the one it displaced.

This is all ahead of the obvious, though nonetheless telling, objection, made by Arthur Fine [4], that any such 'explanationist defence' of scientific realism would be circular. I would put the point as follows. If someone held that it was reasonable to infer to the (approximate) truth of our 'most explanatory' (really best empirically supported) theories, then she would already be a scientific realist—albeit via what I think is the correct route through the union of a series of inferences about particular successful theories, rather than as the result of some fancy meta-level abduction about 'science in general'. For then she would have no need to cite any such meta-level abduction. On the other hand if someone were to question 'inference to the best explanation' (inference to the best supported theory) as regards particular scientific theories, then in consistency she could not fail to question this alleged grand meta-level 'inference to the best explanation' which *at best* (that is laying aside the above objections) has the same logical form.[6]

In sum, then, it surely always was a mistake to think of any 'wholesale' version of scientific realism as a sort of general inference to the best explanation. What are successful or not, what elicit the no miracles intuition or not, are particular *individual* theories—such as Fresnel's wave theory of light or Quantum Field Theory. In so far as there is any sort of 'wholesale' case to be made for scientific realism it is simply as the union of a whole set of specific cases for individual theories.[7] Moreover any

---

[6] Psillos in particular has attempted to defuse Fine's argument; but the attempt to argue in effect that some circles are unproblematic seems to me deeply unconvincing. (See my [22] response to [12] which essays a structurally identical argument.)

[7] There is a fundamental incoherence in the Magnus and Callender paper [14]. While dismissing the 'wholesale' argument for realism, partly on the grounds that the NMA is fallacious, they applaud investigation of 'retail' realist arguments—for particular theories or particular entities (such as 'the' atom). The problem is, of course, that the 'retail' arguments, as suggested above, can all be construed (and ultimately only construed) as instances of some form or other of the NMA. We believe in atoms, because atomic theories have had striking empirical successes to an

such case is made for a *philosophical* thesis not in any sense a 'scientific' one, since there is, and can be, no question of that thesis itself scoring any independent predictive empirical success.

4.3.2 'Retail' Realist Arguments: (Objective) Probabilities Won't Help

Any sensible version of the NMA, then, will be of the 'retail' variety in that its conclusion will be that it is reasonable to hold that some *particular* theory—the wave theory of light, GTR, QFT, ...—is approximately true. Moreover the success involved in the premise of any sensible version will not be any vague, generic, 'wholesale' notion of success but the genuine *predictive* success of the particular theory at issue: the theory must make a prediction of a general kind of empirical result, one that corresponds to the outcomes of observations or experiments. 'Prediction' here, as I have explained elsewhere (see in particular my [23]), need not involve novel, that is, hitherto undiscovered phenomena. The operative condition is that the general phenomenon must not have been 'used in the construction' of the theory at issue (obviously this will automatically be satisfied by any piece of 'new' evidence that was unsuspected at the time when the theory concerned was first formulated). No one is going to exclaim when confronted, say, with some version of Ptolemaic geocentric theory that correctly entails that the planets exhibit stations and retrogressions 'Wow! That must mean that there is something about the theory's fundamental claims that must be at least approximately correct, otherwise it would be a miracle if it succeeded with such a striking prediction'. This is because there is a much more homely explanation of its 'success': parameters in the general Ptolemaic theory (relating sizes of epicycles and deferents, and the relative epicyclic and deferential velocities) had been fixed precisely on the basis of the previous observation of planetary stations and retrogressions, so that the particular version of Ptolemaic theory with parameters fixed in this way was *bound* to yield the phenomena at issue, irrespective of whether or not the overall theory of which it is a part has 'latched on to reality'. This demanding predictivist criterion of success rules out every theory in Laudan's 'plethora' of 'successful' theories that we allegedly now take to be radically false—with one exception: the 'classical' wave theory of light as a periodic disturbance in an elastic medium. Other theories on the list—such as the already mentioned gravitational and physiological ethers of Hartley and Lesage or the astronomical theory of the crystalline spheres—are surely classic instances of *ad hoc* theories. They identify an 'explanatory need'—how, for example, do the sun, planets and stars all move around the earth and why do they all orbit it in the same direction ?—but they 'solve' it (in the geocentric version of crystalline sphere theory by assuming that those astronomical objects are all embedded in concentric spheres

extent that seems entirely implausible if they are not 'on the right lines'. Maybe fancy ways of dressing up 'retail' realist arguments may disguise this fact, but it is a fact nonetheless.

that are themselves revolving in the same direction but at different rates about an axis passing through the Earth) without the slightest hint of any independent testability. The fact that a theory was taken seriously even by serious scientists is not something on which any sensible realist would rest any part of her case. Only predictive success counts.

At least at first blush, the impact of successful prediction for a specific individual theory T can be captured by the following informal argument. T has scored some spectacular predictive success; it would be a miracle if T could get such a phenomenon so exactly right if it were not itself at least approximately correct; but we should not accept that miracles have occurred if there is an alternative explanation of the state of affairs at issue; and there is exactly such a non-miraculous alternative in such cases—namely that T is at least approximately correct; hence we should infer that T is indeed approximately correct.[8]

If we are to capture this argument, and in particular the tricky notion of its being 'a miracle if T were to get evidence e correct without itself being "approximately" or "essentially" correct', in some more formal way, then surely the only realistic prospect is through a probabilistic reconstruction.[9] In investigating the prospects for such a reconstruction, let's first temporarily lay to one side the issues about approximation. The rather nebulous talk about it being a miracle if T had got such a phenomenon as e right if it were not true seems then to translate crisply into the assertion that the probability that e would happen were T false is extremely small: $p(e/\neg T) \approx 0$. And the fact that T (when taken together with accepted auxiliaries) deductively entails e 'translates' of course into the claim that $p(e/T) = 1$. Hence we have:

Premise 1 $p(e/T) = 1$ (e is entailed by T)
Premise 2 $p(e/\neg T) \approx 0$ (it would 'be a miracle if e had been the case were T not true')
Conclusion $p(T/e) \approx 1$ and hence, given that e has occurred, $p(T) \approx 1$.

There are, of course, entirely legitimate worries about what exactly the probabilities in these formulas mean, but laying these worries aside too for the moment (they will

[8] Although the claim that the approximate truth of T would explain its 'otherwise miraculous' success with some surprising prediction e sounds very plausible, it is by no means as obviously true as it might sound. Clearly if a theory is *true* then so are all its consequences—so if it entails some unlikely prediction that turns out to be correct, it seems reasonable to regard the theory's truth as the explanation of its success. But who has shown that all consequences of an 'approximately true' theory (or even, more restrictedly, all the empirically-checkable consequences of such are theory) must themselves be approximately true? Of course, if, as I recommend, 'approximately true' is taken, in structural realist manner, to amount to no more than 'will be retained, modulo the correspondence principle, in all further scientific theories' then this guarantee is supplied.

[9] The other alternative would be to construe the argument as some sort of (allegedly) formal 'Inference to the Best Explanation' which was not itself given a probabilistic construal. My reasons for rejecting this alternative are adumbrated later.

be re-raised very shortly), it is not difficult to show that, so long as they are indeed probabilities, then this reasoning is, as it stands, straightforwardly fallacious.

Here is a simple, and by now well-known, counterexample cited by Colin Howson [6, pp. 52–54]. Suppose that we have a diagnostic test for some disease D, and that this test (unfeasibly) has a zero rate of 'false negatives': that is, the probability of someone's testing negative if she does have the disease is equal to 0; and moreover a non-zero but (again unfeasibly) low 'false positive' rate: say, 1 in a 1000—that is, the probability of someone's testing positive even though they do not in fact have the disease is 1/1000. Suppose now that some particular person x has tested positive, what is the chance that x actually has the disease? In order to avoid changing terminology later, let T stand for the theory that x is suffering from D, while e stands for the evidential statement that x has produced a positive result in the diagnostic test at issue. The zero false negative rate is then just expressed by $p(e/T) = 1$; the low false positive rate by $p(e/\neg T) = 1/1000$; and the probability we are interested in, the probability of x's having the disease given that she has tested positive, is of course $p(T/e)$.

It is often asserted as an empirical result about human psychology (see, for example, [8] and [14]) that most people in these circumstances are inclined to infer from the fact that some person has tested positive and the fact that there is very little chance that x will test positive if she does not have D, that it is highly probable that she does have the disease. Such people would be reasoning in perfect agreement the above version of the NMA:

Premise 1 holds in the diagnostic case because x is certain to test positive (e) if she has the disease (T) (i.e. $p(e/T) = 1$); Premise 2 holds because it is extremely unlikely that x would test positive if she did not have the disease ($p(e/\neg T) = 1/1000 \approx 0$); and the conclusion being drawn is that the probability of x having the disease in view of the positive result—that is, $p(T/e)$—is very high.

Yet, as aficionados are well aware, this inference is an instance of the 'base rate fallacy'. Far from it following that the probability of T given e is very high, *any* non-extreme probability of T, given e, is in fact compatible with the truth of the two premises—even one that is arbitrarily close to zero. It all depends, of course, on the prior probability of T. In the diagnostic case we can, it seems, take that to mean the overall incidence of the disease. If the disease is very rare, a lot rarer than the rate of false positives, then the probability that x has the disease may be very low. So for example, if $p(T) = 10^{-6}$ then the probability here that the person who tested positive has the disease is, via a straightforward application of Bayes's theorem, only around 10^{-3}.

So our first stab at a probabilistic reconstruction of 'the' NMA produces a fallacy. Moreover, the prospects of producing a non-fallacious argument along these lines are surely not improved by reintroducing considerations of approximate, as opposed to 'outright', truth. As we saw, no sensible realist will want to claim anything stronger than that some theory T is approximately true, no matter how astounding its predictive success might have been. But modifying the claim in this way is not likely to help when it comes to reconstructing the NMA probabilistically.

Let $A(T)$ be the assertion that T is approximately true. The relationship between $A(T)$ and e is altogether less clear-cut than that between T and e. I am taking it that, the relevant auxiliaries being taken as given, T deductively entails e; but, on the other hand, $A(T)$, whatever it might precisely mean, presumably need not actually entail e. Nonetheless, since the aim of any version of the NMA is to have e have large impact on $A(T)$—to be reflected, if this reconstruction is to succeed, in an increase in $A(T)$'s probability once e has been observed— presumably the realist will need to claim that $p(e/A(T)) \approx 1$. And again the fundamental assumption here is that the evidence at issue would be very improbable were T not even approximately true, so the realist is presumably committed to the premise $p(e/\neg A(T)) \approx 0$. Hence we have a simple modification of the earlier argument:

Premise 1$'$ $p(e/A(T)) \approx 1$
Premise 2$'$ $p(e/\neg A(T)) \approx 0$
Conclusion $p(A(T)/e) \approx 1$; and hence, given that e has occurred, $p(A(T)) \approx 1$.

But then clearly the base rate problem kicks in just as before: depending on the value of the prior probability of $A(T)$ (the assertion that T is approximately true), any posterior for $A(T)$—including one as close to zero as you like—is compatible with the truth of premises 1$'$ and 2$'$.

If either of the above is the only or uniquely sensible way of capturing the intuitions underlying the NMA, then those intuitions must of course be abandoned entirely since there is no denying the fallaciousness of the base rate fallacy. It seems to me, however, not only that far from being the uniquely correct way to capture those intuitions, it should have been clear ahead of any analysis that no such reconstruction would work. The chief difficulty lies in the issues of how the relevant probabilities could possibly be interpreted in the case of Fresnel's theory or QFT or any other theory whose predictive success elicits the 'no miracles intuition'.

In the diagnostic case, the probabilities involved can arguably be interpreted as objective chances, reflecting —or perhaps constituted by—limiting relative frequencies: the test's false positive rate of 1 in 1000 reflects the assumption that *if* random selections from the whole population were continually made and the frequency recorded of those people who tested positive but failed to have the disease amongst all those testing positive, *then* that frequency would converge on 1/1000 as the number of selections increased indefinitely. Similarly the 'natural prior' in the diagnostic case is the overall incidence of the disease within the population: the proportion of those suffering from the disease is 1 in every million of population and hence if a series of selections were made at random from the population and the relative frequency of those having the disease recorded, then that frequency would converge on 1/1000000 as the number of selections increased indefinitely.[10]

[10] Notice, however, that this is hardly the prior that would 'naturally' be assumed by the Harvard Medical School Students, upon whom much implicit scorn has been poured [6, pp. 52–54]. The fame of this particular case is based on the fact that a (small) group of students at Harvard Medical School allegedly systematically got the 'wrong' answer when asked what the probability is that x

But how should we interpret the probabilities involved in the above probabilistic reconstructions of 'the' NMA—in particular (a) the probability that evidence e would not occur if theory T were false (or not even approximately true), and (b) the 'prior' probability that T is true (or approximately true)? Any attempt to model these probabilities along the lines of those in the diagnostic case would surely be misguided from the start. In order to develop such a model, we would have to think of ourselves as drawing a theory at random from some population of theories and noting whether it was true,[11] how probable it made e and so on.

Notice then that even if we intend to be 'retailist' about the NMA and concentrate on particular successes for particular theories, any attempted probabilistic reconstruction of the argument along these lines forces us back toward at least a somewhat wholesale view: there has to be some reference class of theories, from which the particular theory is regarded as having been drawn and whose characteristics will play an essential role in the argument. But what population of theories, what reference class, should it be?

Despite the intrusion of some wholesale element, it is surely sensible to minimise that element so far as possible. Certainly, then, this reference class of theories should not be thought of as consisting of 'every possible theory' (of what?) —in part because we have no real grasp on what that might be; and also because, in assessing the impact on, say, Fresnel's theory of light of its success with the white spot, there is clearly no interest at all in the fact that theories from, say, chemistry or biology or even other branches of physics fail to entail that same experimental result (why should they?). Moreover, and in line with my criticism of the wholesale approach, neither is there any interest in how many theories from those other scientific fields are true and/or 'successful' in some generic sense.

A more sensible suggestion seems to be that the reference class should consist of rivals to the specific theory for whose likely approximate truth we are arguing. But how liberal should we be with what we count as a rival? It is well known (and strongly emphasised by Howson [6]) that if we count 'gruesome' alternatives, or, in the case of mathematically expressed theories, Jeffreys-style alternatives [12] as rivals, then that class of alternatives will be infinite, indeed non-denumerable. Moreover it

---

has the disease, given that x tested positive (using similar probabilities to those given above). But one assumption involved in the claim that they got the answer about the posterior 'wrong' is that the 'true' base rate that they 'ignored' is the population incidence of the disease. However, no clinician would intuitively 'model' the event of someone's coming through her clinic door as representing a random selection from the population. People don't attend clinics for no reason—the very fact that they are there means that the reasonable guess about the pre-test probability that they have some disease relevant to the clinician's speciality is considerably higher than the population prior. Even in US medicine, where over-testing is rife, the appropriate prior that a patient has some disease ahead of her being subjected to some test, is—thankfully—seldom, if ever, the overall population prior. (For an antidote to the over-investigation venom see [5].)

[11] Of course truth is not an effective notion and so there are bound to be difficulties here too.

[12] Suppose our theory T links two variables and is of the simple form $y = f(x)$; it predicts that when x takes the value x_0, y will take the value y_0; while when $x = x_1$, $y = y_1$; these predictions turn out to be correct when observations are made; Jeffreys pointed out that there are indefinitely many alternatives T' which share this predictive success (at least in the sense that they equally well

is equally well known that major, surely in fact insuperable, difficulties face any attempt to argue that there is an objectively correct prior probability that a theory drawn from such a set of alternatives has some particular property—say (approximate) truth. As for the other crucial probability in probabilistic formulations of the NMA, namely $p(e/\neg T)$ (or, still worse, $p(e/\neg A(T))$), we might start to think of it as measured by the ratio of all possible alternatives to T (or, still murkier, all possible alternatives to $A(T)$) in which e holds compared to all such alternatives. But aside from the fact that we again have no real grasp on what the set of alternatives is, the standard Laplacian chances approach here—as Colin Howson points out [6, p. 46] —is crucially dependent on the assumption that all the basic alternatives are of equal initial weight, and that is surely preposterous in this case.

So in order to arrive at a sensible probabilistic construal of the argument, we would need to restrict in some way the class of rivals to T (or to $A(T)$) that count as part of the appropriate reference class. But how exactly and with what justification? If we restrict the class of alternatives to T's active rivals at the time of its predictive success, this will normally consist of just one theory T' (the corpuscular as opposed to the wave theory of light, classical as opposed to relativistic physics, etc) and $p(e/\neg T)$ is then identified as $p(e/T')$. In the most straightforward case, where we take the theory T' to come along with all the relevant (currently) accepted auxiliaries, then T' will standardly deductively entail $\neg e$. Thus the corpuscular theory of light with natural auxiliaries entails that there will be no 'white spot', classical physics, again with natural auxiliaries, entails an incorrect motion of Mercury's perihelion, which is however correctly accounted for by relativity theory etc. It is easy then to show that the probabilistic version of the NMA goes through without fallacy, since $p(T/e) = 1$. The argument just becomes the probabilistic version of the deductive rule of disjunctive syllogism (and corresponds in the diagnostic case to there being no false positives, which of course means that any person who in fact tests positive must have the disease, *irrespective of base rates*).

But the term $p(e/\neg T)$ in the probabilistic versions of the NMA cannot in fact simply be identified with $p(e/T')$ where T' is T's main historical rival (if, that is, the reconstruction is to capture the underlying intuitions). The possibility that haunts all versions of the NMA is not that some already available theory, different from T, might share the predictive success e at issue—this will demonstrably not be the case.[13] Instead the worry is that some other, *so far unarticulated*, theory could also predict e, while being radically different from T. No one would claim that it was a 'miracle' that T would get some prediction right if it were false, in cases where some known rival T' (that is, a theory that entails that T is indeed false) also made the same prediction. But suppose that T's success is unique—no other *available* theory shares that success. The worry for the realist is arguably that T's

entail the data points (x_0, y_0) and (x_1, y_1)): just take T' as $y = f(x) + (x - x_0)(x - x_1)g(x)$ for any non-zero function $g(x)$. (For more details see [6, pp. 40–44].)

[13] This of course presupposes that the alternative is taken with its 'natural' auxiliaries; the whole basis of the Duhem problem is that the rival can always be made to entail e if we are allowed to add to it any auxiliaries that we like.

success only *seems* ‘otherwise miraculous’ to us precisely because we are unaware of some so-far unarticulated possibility T' that equally well enjoys that predictive success, perhaps has other epistemic virtues, and yet entails that T is way off-beam in terms of what it says is going on at the ‘noumenal’ level. The fact that this so far unknown T' achieves these feats would—however things may seem to us—entail that it would objectively be no miracle for T to have this predictive success despite being false.[14] Or at least T’s success would fail to be ‘miraculous’ in any sense that should incline us to think it likely to be true. The ‘explanation’ in this case would presumably just be that T happens to have the same consequence in respect of e as does the—let’s suppose—true theory T', despite the fact that T is (we are supposing radically) false as revealed by its clash with the true T'. This would be another kind of ‘miracle’ if you like, but one entirely compatible with (indeed one predicated on) T’s falsity.

It seems, then, that if we try think of probabilities like $p(e/\neg T)$ as expressing the ratio of possible alternatives to T in which e holds to all such possible alternatives, then we get into trouble because we have no real handle on that class and certainly no reason to think that all possible alternatives have initial equal weight; but if we restrict the possible alternatives to those we know about (which we might plausibly think about as roughly equal in initial weight), then we also get into trouble since we get trivial answers that have nothing to do with the real issues addressed by the NMA.[15]

4.3.3 The Correct Way to Think About the NMA: The Importance of Not Expecting Too Much

The only serious conclusion to be drawn from the preceding sub-section, so it seems to me, is that there is no available formal probabilistic reconstruction of the NMA that is in anyway convincing because there never was any prospect of producing such a reconstruction. The other proposed reconstruction of ‘retail’ applications of the NMA to particular successful scientific theories involves interpreting them as ‘inferences to the best explanation’. As will perhaps already be clear, I cannot see this as adding anything (except perhaps some confusion) to the intuitions. There *is no method* of inference to the best explanation in any recognisable sense of the word ‘method’. Instead scientists develop theories in various ways, some of these turn out

[14] This is ‘the problem of unconceived alternatives’ mentioned by van Fraassen and given centre stage in a recent book by Kyle Stanford [19].

[15] The situation is clearly not likely to be improved by resort to some intermediate position concerning the relevant ‘population’ of theories—as do Magnus and Callender [14] in identifying this with the class of ‘all candidate theories’. Again this set is ill-defined; again it is hardly likely that each candidate theory will sensibly be modelled as carrying the same weight (or plausibility); and again why should the ratio of successful ‘candidate’ theories that are true (as if we could ascertain this!) in distant fields such as biology or physiology, say, be at all relevant when assessing the impact of the white spot success on the realist credentials of Fresnel’s theory?

to be strikingly successful and are ‘accepted’ as the best available, best empirically supported theories. The suggestion that we are entitled to infer the approximate truth of those theories since it would otherwise be very implausible that they could have been as successful as they have been just is the ‘No Miracle intuition’—to think of this as a case of ‘inference to the best explanation’ adds precisely nothing to that intuition.

Of course it is possible that some other reconstruction can be developed, but it is difficult to see from where. Suppose that the realist in fact concedes that all that she has is the intuition and that she sees no way of a producing a convincing formalisation of the intuitive ‘argument’ linking striking predictive success to truth. This is indeed a ‘concession’ that, again following Poincaré, I was always ready to make.[16] It clearly means that the support for scientific realism (and hence for the structural version I advocate) is modest. But is it entirely non-existent?

That conclusion should be resisted. We should, it seems to me, not expect too much from arguments in philosophy, especially at such a fundamental level as this. There is, of course, no question of a theory's predictive success—no matter how startling and impressive—*proving* that that theory is true (or even ‘approximately true’) and hence solving the problem of (‘vertical’) induction (or ‘abduction’, if you like) at stroke! Perhaps William Whewell believed so. He claimed that the predictive successes enjoyed by the wave theory of light were ‘beyond the power of falsity to counterfeit’. But of course they are not *provably* beyond the power of falsity to counterfeit: the truth *may* be something radically different from what any current theory says it is, and it goes without saying that the (complete) true theory will have all the right empirical consequences, including those describing the predicted effect at issue.[17]

Can we expect to show that, although it is of course possible that the truth is very different from what our current theories say it is, this is at least extremely improbable in the light of their predictive success? Well again surely not in any objective sense of probability—the process of theory production and evaluation, as I have argued, just cannot plausibly be modelled as involving the drawing of theories at random from some super-urn of ‘all possible theories’, or even of all possible rivals to some given theory. We have seen why in some detail in the previous subsection, but I think it ought, on reflection, to have gone without saying.

Proofs and objective probabilities are not what ‘the NMA’ is about. The impact of predictive success, together with the notion of ‘approximate truth’, is ineliminably intuitive—it is of course *possible* that our current theories are radically false despite their predictive success, but this seems so downright *implausible*. Implausible enough to set realism as the default position. It is surely on reflection

[16] In my [21] I refer to the No Miracles ‘consideration’, allowing that it is a mistake to regard it as much of an argument.

[17] Though even Whewell can, I think, plausibly be interpreted as holding only that this is not a ‘realistic’ (as opposed to a merely logical) possibility. Of course we know he was wrong since both Maxwell's theory and photon theory also enjoy the successes at issue and both entail that the classical wave theory is false. But that takes us into the realm of the pessimistic induction.

not surprising that the implausibility here cannot be captured by any sensible analysis in terms of objective probabilities. Realists may wish for something stronger from ‘the’ NMA , but nothing stronger is defensible. It is salutary here to remember Poincaré’s surely correct claim that the NMA intuition is involved, not just in the argument for realism about our scientific theories, but also implicitly in ‘ordinary inductive generalisation’. We have learned to expect that there is no solution of ‘the problem of induction’ (in the original Humean form) either in the form of a convincing deductive argument (by the definition of deductive validity this is *bound* to prejudge the issue!) or in the form of a correct probabilistic argument leading to the conclusion that the generalisation at issue is objectively highly probable, given all the instances. Nonetheless we do not doubt that the reasonable, default, view is that the observational generalisations sanctioned in mature science are in fact correct (though notice that, as Poincaré’s case of Kepler exemplifies, this ‘correctness’ needs to allow for the generalisation’s turning out to be strictly false but still ‘correct within certain limits’). Similarly in the case of the acceptance of ‘observation- transcendent’ theories, which Poincaré—again surely rightly— regarded as simply part and parcel of the same process: the fact that we have no proof and no argument for high objective probability does not imply that, again in appropriate circumstances, the reasonable default position is anything other than that those theories are at any rate approximately correct (a position which also allows—as of course does structural realism—for those theories to turn out to be strictly false but still ‘correct within certain limits’).

So I want to claim that the No Miracles intuition does no more, though also no less, than set some sort of realism as the default position and that it needs no more formal representation in order to do so. Like all arguments for ‘default positions’, the ‘argument’ from some theory T’s predictive success to its approximate truth is defeasible. And indeed it would clearly be defeated *either* by a demonstration that rival theories sharing T’s predictive success but entailing that T is ‘radically’ false can readily and automatically be created; *or* by the demonstration that there are indeed lots of theories from the history of science that were genuinely predictively successful but which can, by no stretch of the imagination, still be seen as ‘approximately true’.

Is there, as some have argued, a demonstration of the automatic availability of ‘equally good’ rivals to accepted theories? Well, as noted earlier, there certainly are well-known constructions that provide alternatives to any given observational generalisation (grue-style constructions) or to any given mathematically expressed theory (Jeffreys-style constructions)—alternatives that share the same empirical consequences as are taken to support the initially given generalisation or theory. But is the fact that these alternatives, by construction, share the same established empirical consequences as their originals enough to establish that they are ‘equally as good’ as those originals? Notice that Poincaré, in the passage quoted from Kepler, talks of its being an unacceptably remote coincidence if all of Tycho’s observations had the planets agreeing with Kepler’s *simple* law and yet—just when neither Tycho nor anyone else was looking— they deviate from their elliptical paths. (‘[I]f a *simple* law has been observed in several particular cases, we may legitimately suppose that it will be true in analogous cases. To refuse to admit this would be to attribute an inadmissible role to chance.') The gruefied or Jeffreys-style theories by construction are all ad hoc and hence do not have the simplicity (or more properly unity) that Poincaré required before they are taken seriously. On my account of confirmation [23], although they have the same consequences as the originals, the gruefied or Jeffreys-style alternatives, unlike the original, gain no empirical support from the phenomena that those consequences describe. This is because both constructions involve parameters that are fixed entirely on the basis of those phenomena. Hence those constructed theories are not in fact 'equally good' as their originals. Admittedly an intuitive judgement lies hidden in this account. This is especially clear in the 'grue' case, since, as everyone knows, if we take grue and bleen as our primitive predicates, then it is the 'all emeralds are green' hypothesis that requires specification of the time parameter on the basis of the observations. We just do need to take for granted some intuitions about which theories in which languages are simple or unified. But again: this is philosophy, we should not expect any more.

Bayesians might seem to supply more, but the appearance is illusory. Bayesians can of course endorse the judgement that the grueified and Jeffreys-style constructions fail to count as 'equally good' as the originals out of which they are created. They can do this simply by pointing out that this will automatically be so if these constructions have considerably smaller prior probabilities than the originals.[18] Similarly Bayesians can endorse Poincaré's account of induction by translating his claim that Kepler's first law is simple into an attribution of reasonably high prior probability to it. In general, as Colin Howson emphasises, there is no problem in supplying a Bayesian reconstruction of the NMA, once any attempt to 'objectify' the argument has been abandoned (for the reasons rehearsed earlier). The fallacy that Howson and following him Magnus and Callender exhibit in probabilistic reconstructions of the NMA is obviously blocked if, far from ignoring the 'base rate', a further premise is incorporated into the argument: a premise that asserts that the prior probability of the theory concerned is not low, but in fact reasonably high.

But of course this Bayesian analysis neither eliminates nor in any sense explains the intuitive judgments involved *either* in the counter to the grue/Jeffreys constructions *or* in the NMA. This is because the evaluation of the prior probability is, on the personalist Bayesian approach advocated by Howson, simply a reflection of a personal judgement about the plausibility of the theory. This means that the Bayesian account is certainly not an improvement on Poincaré's and indeed it seems to me a step backward. The sort of judgment of simplicity or unity that Poincaré pointed to, while it may well be 'subjective' in the sense that it is a judgment that scientists apply without being able to explicate it in more basic terms (and certainly not in terms of objective probabilities), is nonetheless universal within science. It seems to be part of science's very ground-rules that theories with parameters adjusted ad hoc to fit the facts are dispreferred to theories that yield the same facts without the

[18] See [7, chapter 7].

resort to such adhoccery. It seems then to be a mistake to regard this as a personal judgment which an individual ‘agent’ is free to endorse or reject as she sees fit.

Re-focussing on the issue of whether the realist default is defeated by the ever present possibility of constructing empirically equivalent rivals: once this sort of unity or non-*ad hocness* (whether or not regarded as underpinning a high Bayesian prior) is required, then any suggestion evaporates that there are automatic guaranteed ways of generating ‘equally good’ rivals to accepted theories that entail the ‘radical’ falsity of those accepted theories.[19] The remaining threat to the realist default is then the more down to earth or constructive one based on the history of theory change. The worry is that the realist position is defeated by the existence of a long list of actual theories from the history of science that were predictively successful but that cannot any longer sensibly be regarded as even approximately true. I deal with this issue directly in a separate [26] paper. Notice however that *if* this worry can be laid to rest by showing that there is a genuine, if sophisticated, sense in which, despite the considerations raised by history of theory change, the development of science has in fact been ‘essentially’ cumulative, then the default set by ‘the’ NMA becomes stronger. If whenever a theory, despite its predictive success, is eventually replaced, it is invariably replaced by a theory that not only enjoys still further predictive success but substantially retains its predecessor, then the idea that it is very unlikely that our theories fail to be on substantially the right lines surely becomes still more plausible.

This is of course exactly what Structural Realism claims; and it claims that the ‘substantial retention’ occurs at the level of structure.

4.4 Conclusion

In this paper, I have argued only that, despite being used as the starting point for a number of more precise arguments that should never have been taken seriously, the facts about the startling predictive success of some of our theories and the intuitive judgments they elicit still count for something. They provide the very modest basis for a very modest realism. No one should claim that realism can be established in any sense, but the success of some of our theories still seems to make realism about them the most plausible default position. Whether, despite the difficulties that have been raised, structural realism can continue to be defended as a position that not only fails to make the success of our theories a gigantic coincidence , but also, far from being defeated by the facts about theory-change in science, gains support from them, is the subject of forthcoming papers [25, 26]. This is, *contra* Magnus and Callender, not a question that should fill any philosopher of science with ‘ennui’!

[19] See also my [24].

References

1. Boyd, R. (1984) The Current Status of the Scientific Realism Debate. In J. Leplin (ed.), *Scientific Realism* (pp. 41–82). Berkeley, CA: University of California Press.
2. Demopoulos, W. and M. Friedman (1985) Critical Notice: Bertrand Russell's The Analysis of Matter: Its Historical Context and Contemporary Interest. *Philosophy of Science* **52**: 621–639.
3. Duhem, P. (1906) *The Aim and Structure of Physical Theory*, P. Wiener (trans.), Princeton, NJ: Princeton University Press, 1954.
4. Fine, A. (1984) The Natural Ontological Attitude. In J. Leplin (ed.), *Scientific Realism* (pp. 83–107). Berkeley, CA: University of California Press.
5. Gigerenzer, G. (2002) *Reckoning with Risk: Learning to Live with Uncertainty*. London: Penguin.
6. Howson, C. (2000) *Hume's Problem*. Oxford: Oxford University Press.
7. Howson, C. and P.M. Urbach (2007) *Scientific Reasoning: The Bayesian Approach*, 3rd edn. Chicago and La Salle, II: Open Court.
8. Kahneman, D. and A. Tversky (1972) Subjective Probability: A Judgement of Representativeness. *Cognitive Psychology* **3**: 430–454.
9. Ketland, J. (2004) Empirical Adequacy and Ramsification. *British Journal for the Philosophy of Science* **55**: 287–300.
10. Laudan, L. (1981) A Confutation of Convergent Realism. In D. Papineau (ed.), *The Philosophy of Science.* (pp. 139–165) Oxford: Oxford University Press, 1996.
11. Lewis, P. (2001) Why the Pessimistic Induction is a Fallacy. *Synthèse* **12**: 371–380.
12. Lipton, P. (2000) Tracking Track Records. *Proceedings of the Aristotelian Society.* Supplementary Volume **LXXIV**: 179–205.
13. Lipton, P. (2004) *Inference to the Best Explanation*, 2nd edn. London: Routledge.
14. Magnus P. D. and C. Callender (2004) Realist Ennui and the Base Rate Fallacy. *Philosophy of Science* **71**: 320–338.
15. Newman, M.H.A. (1928) Mr. Russell's Causal Theory of Perception. *Mind* **37**: 137–148.
16. Poincaré, H. (1905). *Science and Hypothesis*. Reprinted, New York: Dover, 1952.
17. Psillos, S. (1999) *Scientific Realism—How Science Tracks Truth*. London and New York: Routledge.
18. Putnam, H. (1978). *Meaning and the Moral Sciences.* Boston MA,: Routledge and Kegan Paul.
19. Stanford, P.K. (2006). *Exceeding our Grasp: Science, History and the Problem of Unconceived Alternatives*. New York and Oxford: Oxford University Press.
20. Worrall, J. (1989) Fresnel, Poisson and the White Spot: The Role of Successful Predictions in the Acceptance of Scientific Theories. In D. Gooding, T. Pinch, and S. Shaffer (eds.), *The Uses of Experiment* (pp. 135–157). Cambridge: Cambridge University Press.
21. ———————- (1989) Structural Realism: The Best of Both Worlds. Reprinted in D. Papineau (ed.), *The Philosophy of Science* (pp. 139–165). Oxford: Oxford University Press, 1996.
22. ———————- (2000) Tracking Track Records: A Response to Lipton. *Proceedings of the Aristotelian Society*. Supplementary Volume **LXXIV**: 206–220.
23. ———————- (2002). New Evidence for Old. In P. Gärdenfors et al. (eds.), *In the Scope of Logic, Methodology and Philosophy of Science* (pp. 191–209). Amsterdam: Kluwer.
24. ———————- (2011) Underdetermination, Realism and Empirical Equivalence. *Synthèse* **180**(2): 157–172.
25. ———————- (forthcoming) Defending Structural Realism, Or: The 'Newman Objection' What Objection?
26. ———————- (forthcoming) Avoiding Pessimism the Structural Realist Way.