

Understanding the selective realist defence against the PMI

Peter Vickers¹

Received: 20 July 2015 / Accepted: 23 March 2016

© The Author(s) 2016. This article is published with open access at Springerlink.com

Abstract One of the popular realist responses to the pessimistic meta-induction (PMI) is the ‘selective’ move, where a realist only commits to the ‘working posits’ of a successful theory, and withholds commitment to ‘idle posits’. Antirealists often criticise selective realists for not being able to articulate exactly what is meant by ‘working’ and/or not being able to identify the working posits except in hindsight. This paper aims to establish two results: (i) sometimes a proposition is, in an important sense, ‘doing work’, and yet does *not* warrant realist commitment, and (ii) the realist will be able to respond to PMI-style historical challenges if she can merely show that certain selected posits do *not* require realist commitment (ignoring the question of which posits *do*). These two results act to significantly adjust the dialectic vis-à-vis PMI-style challenges to selective realism.

1 Introduction

The so-called ‘pessimistic meta-induction’ (PMI) challenge to scientific realism has never really gone away since it was vividly articulated by [Laudan \(1981\)](#). Since 1981 the details have changed, but the overall spirit remains the same. Simply put, the history of science still does pose a problem for scientific realists who want to make some sort of inference from the explanatory and/or predictive success of scientific theories to the approximate truth¹ of scientific theories, or theory parts. And it even poses a problem for the ‘selective’ realists, who make a distinction between the ‘working’ parts of a

¹ I follow the convention of using ‘approximate truth’ as shorthand for ‘approximate truth or truth’.

✉ Peter Vickers
peter.vickers@durham.ac.uk

¹ Department of Philosophy, University of Durham, Durham DH1 3HN, UK

theory, which are said to warrant realist commitment, and the ‘idle’ parts of a theory.² Selective realism was actually designed specifically to overcome the PMI. But there remain historical examples where what very much *seem* to be working posits—in the derivation of a successful novel prediction, say—are definitely not approximately true (whatever your theory of ‘approximate truth’). In fact, some of the examples on Laudan’s original list are still relevant here (caloric, phlogiston, the luminiferous ether). But a number of new examples have recently been introduced to the literature, put forward specifically as challenges to selective realism, including: (i) Kirchhoff’s theory of diffraction (Saatsi and Vickers 2011), (ii) Sommerfeld’s prediction of the hydrogen fine structure (Vickers 2012), (iii) Dirac’s prediction of the positron (Pashby 2012), and (iv) Ptolemaic astronomy (Díez and Carman 2015). Vickers (2013) and Lyons (2016) list a number of other possible examples.

In recent years much of the discussion has concerned how (and whether) the selective realist can define ‘working’, such that (a) the definition is properly motivated (not ad hoc), and (b) the realist can use the definition to rebut the historical challenges.³ The thought, usually, is that the realist *has* to be able to define ‘working’, otherwise the position is empty since the realist can’t tell us what we should be epistemically committed to. For example, Stanford (2006, pp.173–180; 2009, pp.385–387) has complained that if the realist can only identify the working posits *in hindsight* (that is, after we already have a successor theory in hand), then realism is bankrupt, since the whole point was to tell us which parts of current science we should/shouldn’t put our epistemic trust in. However, in much of this literature two distinct realist projects have been conflated: (i) the project of responding to the historical challenge, and (ii) the project of explaining what realists should commit to. These projects are not entirely separate, of course, but one must bear in mind that they are definitely not one and the same project. Here—as elsewhere in philosophy—a defence against a challenge (in this case the historical challenge to realism) is not necessarily a positive argument for the view, nor does it have to be.

In this paper I explain what the realist really has to do to respond to contemporary PMI-style objections, emphasising what the realist does *not* have to do. This is articulated as a defensive strategy for the selective realist in Sect. 2. Section 3 then considers a possible ‘disjunction problem’ which arises in this context. In Sect. 4 I consider the possibility that, even if we can’t prospectively identify what warrants realist commitment in our best contemporary theories, we *can* prospectively identify at least *some* elements which do *not* merit realist commitment. In this way it may be possible to make a prediction concerning what will *not* be retained in future science.

² The ‘working/idle’ terminology is now common, and I will adopt it here in a broad sense to refer to what is common to several different contemporary ‘selective’ realisms. These include Kitcher’s distinction between ‘working posits’ and ‘presuppositional posits’ (1993), Psillos’s ‘*divide et impera*’ distinction between ‘idle’ and ‘essentially contributing constituents’ (1999), Saatsi’s focus on ‘success-fuelling properties’ (2005), and Chakravarty’s ‘semirealism’ distinction between ‘detection properties’ and ‘auxiliary properties’ (2007). What these all share is the idea that only certain parts/aspects of a scientific theory are confirmed by the theory’s successes, and thus merit realist commitment.

³ See Peters (2014) for discussion and references.

2 The selective defence against the PMI

This paper will consider one particularly powerful PMI-style objection to scientific realism, consisting in the claim that there are several examples in the history of science of theories which achieved very significant (predictive) success, but where the working posits are not (all) approximately true. In other words (so the claim goes) there are examples in the history of science of derivations of novel predictions, where at least some of the hypotheses which feature in/fuelled the derivation are definitely not approximately true (on any reasonable definition of ‘approximate truth’). The targets of this challenge are most contemporary scientific realists, including those that advocate a selective realist commitment. Even one such example from the history of science can be a real thorn in the side for most realists, who typically think that novel predictive success is a *very* good indication that the theory’s working posits are approximately true. But if there are several examples (and there seem to be) then this argument also speaks against more cautious realists, who claim that novel predictive success is *quite* a good indication, or *probably* means that the theory’s working posits are approximately true.

Let’s think about how the dialectic works here in a little more detail. Since the antirealist is putting forward a challenge to the realist, there is a significant burden on the antirealist to demonstrate the force of the objection. The antirealist needs to present a case from the history of science, identifying a success which is sufficiently impressive for realist commitment—let’s assume a novel predictive success. Then the antirealist needs to reconstruct the derivation of that prediction, identifying assumptions which (at least apparently) merit realist commitment, given their role in the derivation. And then finally the antirealist needs to show that at least one of those assumptions is not approximately true when compared with current scientific thinking. Now this is a lot to achieve, and so the realist has plenty of options for responses. The main options for the realist are as follows:

- (i) Question whether the success identified is really success *enough* for realist commitment.⁴
- (ii) Question whether the reconstruction of the derivation is fair to the history, or whether it is somehow biased, or just one possible reconstruction.⁵
- (iii) Question whether the *specific* ‘working posits’ identified by the antirealist as not approximately true (a) really do merit realist commitment, and (b) really are not-even-approximately-true in light of current scientific thinking.

It should be noted, of course, that these aren’t just options for the realist when responding to an antirealist objection. They are also criteria the realist herself must consider very carefully when she opts to make a realist commitment to certain scientific assumptions, and they can work against her just as much as they can work for her. For example, when considering these criteria it may turn out that some part of science the realist intuitively really *wanted* to believe, or originally took for granted as ‘getting at the truth’ in fact turns out to be something the realist should *not* believe given her own

⁴ Scientific success comes in degrees, even ‘novel predictive’ success. See e.g. [Fahrbach \(2011\)](#).

⁵ Cf. [Schickore \(2006\)](#) here, who argues that all reconstruction ‘requires a perspective’; there is no such thing as the one, true reconstruction. However, see [Kinzel \(2015\)](#) for a response to Schickore.

realist position. So the criteria are very important vis-à-vis responding to antirealist threats, but they can also force realists to withdraw realist commitment from parts of science they really *want* to commit to (e.g. because, intuitively, there seems to be ‘lots of evidence’ for the scientific claims in question).

The focus of this paper is going to be criterion (iiiia), questioning whether the specific assumptions identified by the antirealist really do merit realist commitment. It is here that much of the literature on this topic (including Psillos 1999; Lyons 2006; Vickers 2012, and Harker 2013) has conflated the realist’s defence against the PMI and the positive project of the realist to identify what the realist should be committed to in a given case. Now, *one* way one *might* respond to an antirealist challenge is to invoke some theory of ‘working posits’—identifying which assumptions really are ‘working’ in the case in question—and then conclude that the posits identified by the antirealist are not included within the working posits. One of the most influential figures in the selective realist community is Stathis Psillos (e.g. Psillos 1999) and he is very easily interpreted as doing just this. Lyons (2006) introduces Psillos’s divide et impera realism as a response to ‘the historical challenge’, shows that there are problems with Psillos’s definition of ‘the posits which really fuel the derivation’ (which merit realist commitment according to Psillos), and concludes (p.537) that “this sophisticated form of realism remains threatened by the historical argument that prompted it.” Harker (2013)—partly drawing on Lyons (2006)—similarly fuses together both the realist’s defence against PMI-style objections and the positive project of the realist. This leads him to state the following of the selective realist strategy:

For the strategy to answer Laudan-style concerns [...] the criteria we invoke to isolate those constituents of theories that are to be recommended for realist endorsement must render such constituents epistemically accessible (Harker 2013, Sect. 2).

This entails that to answer the historical challenge we need to have a convincing theory concerning how it is possible to identify the posits which warrant realist commitment. But this is incorrect. To answer the historical challenge the realist *can* do this, and prominent realists such as Psillos *have* done this, but the realist *need* not do this. Indeed, to develop a theory of which posits warrant realist commitment is to make the realist’s task much harder than it has to be if all she wishes to do (for now) is respond to ‘Laudan-style concerns’.

To respond, all the realist needs to do is show that the specific assumptions identified by the antirealist do *not* merit realist commitment. And she can do this without saying anything about how to identify the posits which *do* merit realist commitment. How is that possible? There is more than one answer here. The simplest case is when an assumption can just be eliminated without affecting the derivation in question. In this case the assumption is clearly idle vis-à-vis the success, and thus doesn’t merit realist commitment. And this can be established without any real theory of what *does* merit realist commitment, simply by eliminating the posit in question and displaying the resultant derivation. However, such a case is going to be vanishingly rare. If an assumption is so obviously idle, then (a) scientists would usually have left it out in the first place, and (b) antirealists looking for a serious threat to selective realism would

surely recognise that the assumption in question is not doing any work, such that the selective realist has an easy response to the challenge.

Much more serious is a case where a scientist has indeed *used* a posit in the derivation of a novel prediction, and that posit *cannot* be simply eliminated from the derivation without also eliminating the success. This means that there is a straight forward sense in which the assumption *is* doing work: the derivation doesn't go through without it. But at the same time the posit in question does *not* merit realist commitment, because it is not confirmed by the success. The reason is this: the posit in question is doing work in the derivation *solely in virtue of the fact that it entails some other proposition, which itself is sufficient (when combined with the other assumptions in play) for that specific derivational step*. In such a case it is this other proposition, loosely speaking 'contained within' the original proposition, which is (seems to be) really fuelling that particular step in the derivation. The bracketed 'seems to be' signals the fact that, again, the realist doesn't need to identify what really deserves realist commitment here. If all the realist wants to do, for now, is respond to the antirealist challenge, then all that matters is that the original proposition does *not* merit realist commitment. And this is the case even though there is an important sense in which it is not completely idle.⁶

We can bring this sort of case to life with some examples, first a couple of toy examples and then a real case from the history of science. Consider first a situation where a doctor supposes you have the adenovirus (e.g. because the adenovirus is known to be widespread in the neighbourhood). That doctor might well use this assumption along with other assumptions about the human immune system to quite accurately predict how your symptoms will develop. But the doctor might be wrong about your having the adenovirus, and in addition the doctor's assumption is doing work in the sense that she reaches her conclusion making use of that assumption. Is this a case of miraculous success, then, since the doctor predicted correctly whilst making use of a false assumption?

Of course not: it's easy to make sense of how the doctor predicted correctly despite her mistake. The truth is you do have *one* of the cold viruses. And the doctor's reasoning only depended on her (implicit) belief that you have *one* of the cold viruses, a belief she is committed to in virtue of the fact that she believes that you have the adenovirus. Thus her false belief is doing work in her reasoning *solely in virtue of the fact that it entails some other proposition, which itself is sufficient (when combined with the other assumptions in play) for the success*. Another way to put it is this: the doctor was committed to what mattered here; she was mistaken only about redundant details, details which go beyond what was needed to make her prediction.

This particular way of identifying a proposition as not meriting realist commitment *has* been touched on in the literature, but it is not widely appreciated. For example, Saatsi (2005, p. 532) discusses a case concerning a crammed elevator which refuses to move. Saatsi is interested in *explanations* which contain false content which is 'surplus' in the sense that it is non-explanatory. But the case is also useful when

⁶ In previous work (e.g. Vickers 2013) I have tended to describe the original proposition as 'idle' in such a case. I now believe that to define 'idle' in the necessary way one is going sufficiently against many intuitions that confusion is bound to result. Better to stick with definitions which are closer to common intuitions, and describe the original proposition as 'doing work'. Thus the working/idle distinction no longer maps onto the merits/doesn't merit realist commitment distinction.

trying to understand realist commitments in a case of predictive success. Suppose somebody predicts that an elevator will refuse to move by reasoning with the false assumption that the elevator load is 50 kg too heavy. One predicts successfully, since the reasoning depends only on the assumption that the load is too heavy—the belief that it is specifically 50 kg too heavy is redundant detail. Or to put it another way, the ‘50 kg too heavy’ assumption does work in the derivation *only in virtue of the fact that it entails another ‘too heavy’ assumption which itself is sufficient, when combined with the other assumptions in play, to reach the true prediction.*

Many in the community still assume that if an assumption is *used* in a derivation then the realist must make a commitment to it (must believe it to be approximately true). Lyons (2006) even describes the selective realist strategy as ‘deployment realism’ (following Kitcher 1993), terminology which strongly suggests that the ‘deployed’ (used) assumptions are the assumptions the realist must commit to. But this is not the case: the given toy examples show very clearly that, at least sometimes, realists *should not* make a realist commitment to all the assumptions employed within a derivation of a successful prediction. After all, it would be madness for anyone to believe the lift to be exactly 50 kg too heavy just because that assumption was *used* to reach a correct prediction.

We can illuminate the strategy further with a real case from the history of science. Take Bohr’s prediction of the frequencies of the spectral lines of ionised helium. Vickers (2012) puts this forward as a possible counterexample to selective realism. It seems to fit the bill, since:

- (a) The prediction of the frequencies of the spectral lines of ionised helium was a novel predictive success, and was seen as extremely significant at the time.
- (b) Bohr’s prediction came about by making direct use of his theory of the atom (which includes some assumptions which are definitely not approximately true).

However, Vickers (2012) presents a ‘way out’ for the selective scientific realist, by noting that one of the assumptions Bohr made, and *used* to reach his successful predictions, does not merit realist commitment. The assumption in question is as follows:

H: The electron orbits the nucleus at specific, quantised energies, corresponding to only certain ‘allowed’ orbital trajectories.

This doesn’t merit realist commitment, since it can be seen to do work within the relevant derivation *solely in virtue of the fact that it entails another proposition, which itself is sufficient for the relevant step of the derivation.* This other proposition is the following:

H*: The electron can only occupy certain, specific, quantised energy states within the atom.

Bohr was committed to H* in virtue of his being committed to H. But it turns out that H* is sufficient for Bohr’s derivation to go through.⁷ And, crucially for the selective

⁷ Cf. Norton (2000, pp. 86–87): “The reduced form [of the derivation] eschews all talk of elliptical orbits other than in the domain of correspondence with classical theory... No assumption is made or needed that these stationary states are elliptical orbits of some definite size and frequency of localised electrons. What is retained is that these states possess a definite energy.” See also Vickers (2012).

realist, whilst H is certainly false (there are no ‘trajectories’), H^* is approximately true (by the lights of current scientific thought).

This is enough to answer the antirealist’s concern that one of the working posits was not approximately true. The realist might well accept that H is not approximately true, given the reference to orbital trajectories. The realist might also accept that H is ‘doing work’ within Bohr’s derivation, in the sense that it certainly cannot be simply eliminated from the derivation without destroying the derivation, and thus the success. However, H does not merit realist commitment, given its relation to the success in question. There is an important sense in which it is not *directly* fuelling that success; rather it is only *indirectly* fuelling that success in the sense that it does so via its relation to H^* . If anything merits realist commitment here, it is H^* , and not H .

The crucial step now is to note that the realist does *not* need to claim that H^* merits realist commitment. All that matters to answering the challenge is to show that H does *not* merit realist commitment (at least, not in virtue of its relation to Bohr’s prediction of the ionised helium spectral lines). Vickers (2012) actually muddies the waters here: he makes a distinction between (i) the assumptions Bohr *used* to reach his predictions, and (ii) the assumptions which were ‘truly necessary’ to generate the predictions (p. 10). But the words ‘truly necessary’ are a mistake—they belong in a discussion of what the realist should be committed to, *not* in a discussion of what the realist should *not* be committed to. And to answer the challenge the realist just doesn’t need to claim that H^* is ‘truly necessary’ for the derivation. Perhaps the derivation can go through with an assumption *still weaker*, such that it turns out that H^* is *also* immune to realist commitment. Perhaps only very abstract ‘structure’ truly merits realist commitment, as structural realists like to claim. But that can be left for another day. We are not here in the business of identifying realist commitments; we are in the business of showing that some specific assumption does *not* merit realist commitment. Because that is enough to answer the historical challenge.

3 A disjunction problem?

The toy examples above serve an important role in this paper. They serve to *show by example* just how ridiculous it would be to suggest that an assumption merits doxastic commitment, on ‘no miracles’ grounds, just because it was *used* to reach a successful prediction. A full answer to the question ‘Why exactly is it ridiculous?’ is harder to provide. Nonetheless, I do have a partial explanation to give: sometimes it is clear that a *used* assumption is not confirmed by the success it leads to because it is so clear that it did work to generate the success *solely in virtue of the fact that it entails some other proposition which itself is sufficient for the derivational step in question*. This is just a partial explication, of course, but hopefully it is explication enough to have persuasive force. Another worry is that, far from requiring further explication, it already fails as it stands because of the use of ‘entailment’ as a key part of the explication.

Entailment seems to work well for the examples given in the previous section: ‘Dave has one of the cold viruses’ is entailed by ‘Dave has the adenovirus’, ‘The load is too heavy’ is entailed by ‘The load is 50 kg too heavy’, and ‘The electron can only occupy certain energy states’ is entailed by ‘The electron can only occupy

certain orbital trajectories'. So why not entailment? Indeed, Vickers (2013, p. 198) considers precisely the sort of example under consideration here, and talks in terms of one proposition (the original proposition used within the derivation) 'containing within it' another proposition which can take the place of the original proposition in the derivation. Vickers (ibid.) then writes "By 'contain within it', the realist means simply that some weaker proposition can be inferred from the original proposition: so 'The passengers are too heavy' is contained within 'The passengers are 50 kg too heavy'." However, Vickers' account leads to some rather awkward results. In particular, it leads to a disjunction problem.

The basic problem is the Principle of Addition: any given proposition P entails $P \vee Q$ for any arbitrary proposition Q. Now, as explained above, to respond to an historical challenge a realist only needs to show that some individual proposition P does *not* merit realist commitment. And Vickers' analysis seems to suggest that if the realist can find *any* proposition entailed by P which can take the place of P in the derivation, then that shows that P does not merit realist commitment.⁸ The worry with this is that P entails $P \vee Q$ for *any* Q whatsoever, and so Q can be selected to *make sure* that $P \vee Q$ can take the place of P in the derivation without affecting the derivation of the prediction in question (call the original prediction 'A'). Since $P \vee Q$ is a disjunction one might worry that any resultant prediction would also be disjunctive, and so couldn't be the same as the original prediction, A, achieved by using P. But this isn't necessarily the case: if we use $P \vee Q$ in place of P the final prediction could take the form $A \vee A$, and thus collapse to the original prediction A. This might well be very difficult to achieve in practice, and also not representative of serious science, or philosophy. But the point is that Vickers' analysis *allows* for this absurd response to an antirealist challenge, and that is unacceptable. It should be ruled out from the start.

One option here is for Vickers (2013) to retain his analysis but reject the Principle of Addition, perhaps by adopting a relevance logic. Certainly one could find support for this move in the literature. Weingartner (1993) notes that "[T]his principle [of addition] is the culprit of a lot of difficulties in different areas," and argues that it is responsible for "most of the well-known paradoxes in the theory of explanation, confirmation, law statements, disposition predicates, etc." His conclusion is that we ought to put limitations on the application of classical logic, especially "if we think of logical consequences drawn in science from assumptions or hypotheses" (p. 95). More recently Strevens (2008, Sect. 3.61) encountered his own 'disjunction problem' in the context of developing a causal theory of explanation. And his ingenious attempts to get around the problem—making reference to 'causal contiguity at the fundamental level'—are controversial (see Strevens et al. 2012 for discussion). Perhaps Strevens too could consider adopting a relevance logic, or (following Weingartner) adopting classical logic but limiting its application. This will look like an *ad hoc* solution to many, but recent work on logical pluralism encourages us to think of logic as a tool,

⁸ What of derivations consisting of a single assumption entailing another single assumption entailing another single assumption and so on? Such derivations would pose a problem for my account. In response, I simply don't take such 'derivations' seriously as a part of real science: real scientific derivations consist of various different assumptions coming together in a series of complex multi-assumption derivational steps.

not a truth (cf. [Beall and Restall 2006](#)). In which case we need not be bound by every rule of classical logic in every conceivable context.

I am sympathetic to this line of argument, but in fact it isn't necessary to mess with the logic. There is a crucial difference between the (brief) analysis given by [Vickers \(2013\)](#) and the analysis given here, such that the noted disjunction problem affects only [Vickers \(2013\)](#). The key to avoiding the problem lies in the particular wording: *P does not merit realist commitment whenever P is doing work in the derivation solely in virtue of the fact that it entails some other proposition which itself is sufficient, when combined with the other assumptions in play, for the relevant derivational step*. To see this, suppose that a given P is swapped in the derivation for $P \vee Q$, for some random proposition Q. Certainly it is the case that P entails $P \vee Q$, and it might be the case (if Q is carefully selected) that $P \vee Q$ itself is sufficient for the relevant derivational step. But what's missing here is the condition that P is doing work in the original derivation *in virtue of the fact that it entails $P \vee Q$* . Remember that Q has been carefully selected to ensure that the final prediction is unaffected. Most probably, it didn't feature at all in the original derivation or indeed the relevant history of science. How, then, could it make sense to say that the work done by P in the original derivation is work done in virtue of the fact that P entails $P \vee Q$? Entailment of $P \vee Q$ from P is not enough to meet this condition: we need something more than mere entailment. One option might be to turn to *entailment which is metaphysically necessary*. But to pursue this further would open up a huge can of worms in metaphysics and logic, e.g. concerning the concept of a *grounding* relation (see e.g. [Raven 2015](#)). All that matters for present purposes is that it does *not* make sense to say that a given proposition P does work in a derivation in virtue of the fact that it entails $P \vee Q$ where Q is selected specifically to leave the final prediction unaffected and did not feature at all in the relevant history of science.

4 Prospective identification of posits *not* meriting realist commitment

As argued in Sect. 2, when a specific historical challenge is put forward, the realist can answer that challenge by arguing that the specific hypotheses in question do not merit realist commitment (staying silent on what *does*). Consider again the case of Bohr's theory of the atom, and his hypothesis concerning electron orbits. Although this hypothesis was *used* by Bohr, the realist can argue that it only did work for Bohr in virtue of the fact that it entails some other proposition which is itself sufficient, when combined with the other assumptions in play, for the relevant derivational step (and thus for the final success). Specifically, Bohr's derivation of the spectral lines of ionised helium goes through just the same if we make reference only to electron *energies*. Now one may wonder whether this could have been noticed *at the time*, or whether we can only reach this conclusion in hindsight. There has been much fuss of the distinction between prospective and retrospective identification of working posits. [Stanford \(2006, pp. 167–180; 2009, pp. 385–387\)](#) has argued that we *cannot* prospectively identify working posits, and that this is a major blow to realism, since this is "what the realist cared most about all along" ([2009, p. 387](#)). But even if Stanford is right (I stay silent on this question) an interesting alternative for the realist is to at least prospectively narrow down the targets of realist commitment, by identifying at

least *some* of the posits which *don't* merit realist commitment even though they are working posits, in the sense noted above. And this is much more feasible, I submit, simply because it is so much easier to identify something a realist should *not* commit to than it is to identify something a realist *should* commit to.

How would we go about prospectively identifying at least some posits which, despite being working posits, are not confirmed by the success such that they merit realist commitment as a consequence of that success? Well, we can simply go through the assumptions in a derivation one at a time, and consider the following question: Is this assumption doing work in the derivation solely in virtue of the fact that it entails some other proposition which is itself sufficient, when combined with the other assumptions in play, for the success?⁹ Consider how this could have worked out in the Bohr case. How might it have been possible to see that his hypotheses concerning electron orbits were not confirmed by the success? Well, Bohr's assumption H concerning quantised electron *trajectories* entails H*, which merely concerns quantised electron *energies*. Bohr was certainly in a position to notice this entailment relation. And he was also, of course, in a position to notice that H* is sufficient for the success: H* can take the place of H in the derivation without affecting the result. Was Bohr *also* in a position to note that H does work in his original derivation solely in virtue of the fact that it entails H*? It would seem so, since the fact that H* is sufficient for the result shows that the reference to trajectories is redundant vis-à-vis the success. If this is right we have the prospective identification of a posit which is not confirmed by the success it leads to, and this might even lead to a prediction concerning the future development of the relevant science, namely, reference to the orbital trajectories of electrons will not be retained in the successor theory. And this prediction of course would have turned out to be true!

We are perhaps getting ahead of ourselves. Putting ourselves in Bohr's shoes for a moment, it may have been inconceivable at the time to think that electrons could have quantised electron energies *without* having associated quantised orbital trajectories (cf. [Stanford 2006](#), p. 171). But this wouldn't have stopped Bohr noticing that his references to electron orbits were redundant vis-à-vis his predictive successes. Or, to put it another way, that his references to electron orbits were not confirmed by the success. Naturally in such circumstances one might still want to believe in quantised electron orbits on the grounds that these are, apparently, metaphysically necessitated by the quantised electron energies and other relevant assumptions. But at least Bohr could have separated two importantly different motivations for his beliefs: his beliefs in the quantised energies were motivated directly by their role in generating the successful predictions, but his beliefs about quantised orbits were motivated *by an inference from his beliefs concerning quantised energies*. And, since his beliefs about energies do

⁹ There is an interesting problem lurking here, which I merely present for future investigation. Suppose a given derivational step consists of two assumptions A and B combining to deliver C. And suppose that A entails A* and A* combined with B still delivers C. In that case A does not merit realist commitment. Suppose further that B entails B*, and B* combined with A still delivers C. In that case, B does not merit realist commitment. But now suppose that A* and B* together do *not* deliver C. In that case it seems that what the realist should *not* be committed to is underdetermined, since she has equal reason not to be committed to A and not to be committed to B. And yet not committing to *both* A and B is not an option given that A* and B* combined do not deliver C.

not properly (logically) *entail* his beliefs concerning orbital trajectories, he could have come to agree that the latter were not as secure as the former. This gives us a way to answer Stanford's question: "[W]hy did we (or the relevant scientific communities) *ever* believe more than those parts or aspects of past theories on which their empirical successes really depended?" (2009, p.385). The answer in Bohr's case is that the assumptions apparently responsible for his empirical successes (concerning electron energies) *appeared to conceptually entail other assumptions* (concerning electron orbits). But it still stands that they weren't directly confirmed by the success, because they can be seen to be redundant vis-à-vis that success.

This all puts us in a position to be optimistic concerning the prospects for prospectively identifying at least some of the posits which do *not* merit realist commitment in our current best scientific theories. One can at least attempt to go through modern-day derivations of predictions of phenomena, and look for hypotheses which include details which are redundant for the purposes of the derivation. In this way it may be possible—at least *sometimes*—to separate the elements of theory which are confirmed by the phenomena they predict, and the elements which are not so-confirmed, but which we believe for other reasons (e.g. because they appear to be conceptually or metaphysically entailed by the elements which *are* confirmed). And when it comes to the bizarre world of fundamental physics we might come to agree that these *other* reasons have been shown time and time again through the history of science not to be good reasons. Thus the realist should, perhaps, restrict her commitments to what is *directly confirmed* by the predictive successes.

At the very least this seems to me to be a worthwhile heuristic to bear in mind when we (scientists in particular) are thinking about how scientific progress might be made. Including how general relativity and quantum theory might ultimately be reconciled. Anything that can help with this extraordinary challenge can only be a good thing.

Acknowledgments This research was undertaken as part of the AHRC project 'Contemporary Scientific Realism and the Challenge from the History of Science', grant reference AH/L011646/1. The author would like to give particular thanks to the audience at the Cape Town conference 'New Thinking About Scientific Realism' (August 2014), where an earlier version of this paper was presented, with especial thanks to Juha Saatsi, Dana Tulodziecki, and David Harker. The author would also like to thank two anonymous referees for probing criticisms.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Beall, J., & Restall, G. (2006). *Logical pluralism*. Oxford: Clarendon Press.
- Chakravartty, A. (2007). *A metaphysics for scientific realism: Knowing the unobservable*. Cambridge: Cambridge University Press.
- Díez, J. A., & Carman, C. (2015). 'Did Ptolemy make novel predictions? Launching Ptolemaic astronomy into the scientific realism debate. *Studies in the History and Philosophy of Science*, 52, 20–34.
- Fahrbach, L. (2011). Theory change and degrees of success. *Philosophy of Science*, 78(5), 1283–1292.
- Harker, D. (2013). How to split a theory: Defending selective realism and convergence without proximity. *British Journal for the Philosophy of Science*, 64(1), 79–106.

- Kitcher, P. (1993). *The advancement of science: Science without legend, objectivity without illusions*. Oxford: Oxford University Press.
- Kinzel, K. (2015). Narrative and evidence: How can case studies from the history of science support claims in the philosophy of science? *Studies in the History and Philosophy of Science*, 49, 48–57.
- Laudan, L. (1981). A confutation of convergent realism. *Philosophy of Science*, 48, 19–48.
- Lyons, T. (2006). Scientific realism and the Stratagema de Divide et Impera. *British Journal for the Philosophy of Science*, 57, 537–560.
- Lyons, T. D. (2016). Structural realism versus deployment realism: a comparative evaluation. In C. Haufe (Ed.), *Studies in the history and philosophy of science*, special issue on Testing structural realism (forthcoming).
- Norton, J. (2000). How we know about electrons. In R. Nola & H. Sankey (Eds.), *After Popper* (pp. 67–97). Kuhn and Feyerabend: Kluwer.
- Pashby, T. (2012). Dirac's prediction of the positron: A case study for the current scientific realism debate. *Perspectives on Science*, 20(4), 440–475.
- Peters, D. (2014). What elements of successful scientific theories are the correct targets for "selective" scientific realism? *Philosophy of Science*, 81(3), 377–397.
- Psillos, S. (1999). *Scientific realism: How science tracks truth*. London; New York: Routledge.
- Psillos, S. (2009). 'Grasping at Realist Straws', review symposium of Stanford (2006). *Metascience*, 18, 363–370.
- Raven, M. J. (2015). Ground. *Philosophy Compass*, 10(5), 322–333.
- Saatsi, J. (2005). Reconsidering the Fresnel–Maxwell case study. *Studies in History and Philosophy of Science*, 36, 509–538.
- Saatsi, J., & Vickers, P. (2011). 'Miraculous Success? Inconsistency and untruth in Kirchhoff's diffraction theory. *British Journal for the Philosophy of Science*, 62(1), 29–46.
- Schickore, J. (2011). More thoughts on HPS: Another 20 years later. *Perspectives on Science*, 19, 453–481.
- Stanford, P. K. (2006). *Exceeding our grasp*. Oxford: OUP.
- Stanford, P. K. (2009). 'Author's Response', in 'Grasping at Realist Straws', a review symposium of Stanford (2006). *Metascience*, 18, 379–390.
- Strevens, M. (2008). *Depth*. Harvard: Harvard University Press.
- Strevens, M., Weatherson, B., Hall, N., & Lange, M. (2012). Book symposium on depth by Michael Strevens. *Philosophy and Phenomenological Research*, LXXXIV(2), 447–505.
- Vickers, P. (2012). Historical magic in old quantum theory? *European Journal for Philosophy of Science*, 2(1), 1–19.
- Vickers, P. (2013). A confrontation of convergent realism. *Philosophy of Science*, 80(2), 189–211.
- Weingartner, P. (1993). Can there be reasons for putting limitations on classical logic? In P. Humphreys (Ed.), *Patrick suppes: Scientific philosopher* (pp. 89–124). Dordrecht: Kluwer Academic Publishers.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: OUP.