

Structural realism: the only defensible realist game in town?

John Worrall

1. *Predictive Success, Theory-Change and Scientific Realism*
 2. *Selective Realism and its Problems*
 3. *Entity Realism*
 4. *Criticisms of Structural Realism - articulated and countered*
 - 4.1 *The Fresnel- Maxwell example is entirely unrepresentative*
 - 4.2 *'The Newman Objection' (or rather 'The Newman "Objection"')*
 5. *Conclusion*
-

1 Predictive Success, Theory-Change and Scientific Realism

Most of us share the 'no miracles intuition': the feeling that the success of our currently accepted scientific theories – particularly their success in predicting hitherto unsuspected phenomena – provides at least some reason to hold that what those theories say that goes beyond the phenomena, and generalisations thereof, is likely to be true. Or, to be a little more reflective and fine-grained about the intuition from the outset: that such predictive success provides some reason for holding that what those theories say is going on 'behind' the phenomena is likely to be at least *in part*, at least *approximately* true.

Of course, intuitions should not be given free rein, but instead be subjected to critical analysis. And in fact, it takes only a moment's reflection to see that no one should be what I call a 'gung ho realist'. A 'gung ho realist' believes that our current theories are so successful that the only reasonable belief about them is that they are true. But in fact, no one should think that both the Standard Model of Quantum Mechanics and the General Theory of Relativity – arguably our best current theories – are likely to be literally true. The Standard Model involves singularities that make a fully realist interpretation impossible; and, although it and General Relativity are, arguably, not formally inconsistent, they are extraordinarily uncomfortable bedfellows. Physicists expect an eventual, though currently unarticulated, 'synthesis' that will involve corrections or emendations of at least one, and more probably both, of the two theories.

So, no one should think that our current basic theories are any more than 'approximately true' (whatever that turns out to mean). Perhaps less obviously, *even ahead of any consideration of the history of theory-change in science*, no one should believe that there are grounds for holding that every part of every successful scientific theory is even approximately true. Consider a couple of illustrative historical examples:

1. Kepler, Galileo and Newton all believed that it was reasonable, in the light of its empirical success, to hold an attenuated realist view of the Copernican theory – to believe that Copernican theory not only 'saves the phenomena' more elegantly than its Ptolemaic rival, but also that what the theory says about the motion of the earth about its own axis and about the sun is 'essentially'

true. But they did not believe that a realist attitude was justified towards Copernicus's claim about a 'third motion' – a supposed conical motion of the earth's polar axis (Kuhn [1957], pp. 164–165). These luminaries all (correctly) held the 'third motion' to be an entirely *ad hoc* element of Copernican theory rendered necessary only because a part of Copernicus's overall theory (his commitment to the idea that the planets are stuck in crystalline spheres whose motion they consequently share) is entirely false. Hence, Kepler, Galileo and Newton show that being a realist about a theory 'as a whole' need not involve being realist about the whole theory!

2. No one should have believed (and many in the eighteenth- and nineteenth-centuries did not believe) that the success of Newton's theory of universal gravitation – in, for example, predicting the precession of the equinoxes, or, later, in predicting the existence of the hitherto unsuspected planet, Neptune – justified a realist attitude toward the whole of that theory, including the assertion (which Newton explicitly labelled an 'hypothesis') that the centre of mass of the universe is at rest in absolute space. This is because, as Newton himself proved, exactly the same phenomena – *all* the phenomena *in principle*, not just those so far observed – follow, not only from his overall theory, *T*, including that 'hypothesis', but just as well from any theory *T'* that differs from *T* by assuming *any* non-zero (though uniform) absolute velocity for the centre of mass of the universe (Newton [1934], Bk. I, pp. 165–166 and Bk. III, pp. 574–575). There is, then, a clear sense in which, ahead of any consideration about the eventual replacement of his whole theory, Newton's assumption of absolute rest was redundant and hence irrelevant for the success of his overall theory. So, there is no realist warrant for taking the success of the theory as a whole as likely to betoken the truth, or even approximate truth, of that particular part of it.

To emphasise: these examples (and others) involve no consideration of overall theory-change – we did not need the vantage point of later, better theories than Copernicus's to identify his assumption of the 'third motion' as *ad hoc*; nor did we need the vantage point of Einstein's replacement of Newton's theory to identify as redundant the latter's hypothesis about the centre of mass of the universe being at absolute rest. The 'no miracles intuition' that motivates realism clearly does not apply in such cases. Would it be a 'miracle' if Newton's theory enjoyed the predictive successes that it does if his assumption ('hypothesis') that the centre of mass of the universe is at absolute rest were wrong? Well, clearly not – Newton himself showed that an infinite range of alternatives (each of which entails, of course, that Newton's favoured hypothesis is indeed false) produces the same predictive success. Equally, because it is not independently testable (it simply corrects what would otherwise be a false observational consequence about the inclination of the earth's axis during its orbit around the sun), Copernicus's postulation of the 'third motion' produces no apparent 'miracles' to be explained away. The sensible scientific realist, then, is automatically, ahead of any consideration of theory-change in the history of science, not only an *approximative* realist but, more interestingly, possibly a partial or *selective* realist (depending on detailed features of the theory involved).

However, when we *do* turn to the facts about theory-change in science, even this more sensible, less 'gung ho' realism might seem to be untenable – at least at first sight.

Hackneyed examples are often hackneyed for good reasons. So, let's again consider the hackneyed example. Fresnel's classical wave theory of light, as developed in his 1818 memoir on diffraction and later work, was notably predictively successful. As has been frequently discussed, it correctly predicted that if a small opaque disk is held in the light diverging from a point source then, far from casting a solid shadow on an observation screen some distance away (as geometrical optics would predict), that 'shadow' has a white spot at its centre which is just as intensely illuminated as if no opaque disk were there. Fresnel's theory also predicted a range of other diffraction and interference results. Moreover, as a consequence of Fresnel's shift in the

early 1820s to the claim that light waves are transverse rather than longitudinal, his theory predicted the existence of the hitherto entirely unsuspected phenomena of internal and external conical refraction. I now accept that I have in the past been guilty – along with other realist philosophers of science – of putting undue weight on the pro-realist power of *single* predictive successes, such as that with the white spot. But, of course, logic tells us that every false theory has true deductive consequences, indeed infinitely many of them. And I now think that the realist should concede that there have been theories – arguable examples are the caloric theory of heat and even Ptolemy’s geostatic theory (see Ladyman [2011] and Carman and Díez [2015]) – that were, so to speak, ‘very false’ – not even approximately true – even though those theories enjoyed novel predictive success. The situations which really underwrite the case for some sort of realism, the situations that really ginger up the ‘no miracles intuition’ are not ones in which a theory enjoys an odd isolated predictive success, but instead ones in which that theory provides a *range* of different novel and empirically correct predictions. Applying this intuition to our hackneyed case: wouldn’t it be amazing if a theory like Fresnel’s could make so many astonishing predictions about novel phenomena, not only interference and diffraction results but also internal and external conical refraction, and yet still be radically false in what it says is going on at the ‘noumenal’ level to produce those phenomena?

The well-known rub, however, is that what Fresnel’s theory says is going on at the ‘noumenal’ level, ‘beneath’ the phenomena is this: vibrations created in light sources, such as the sun, are transmitted to, and then through, an all-pervading elastic medium – the initial vibrations in the source force the neighbouring particles of the medium away from their positions of equilibrium, whereupon elastic restoring forces bring those particles back toward their equilibrium positions creating the disturbances that spread to neighbouring particles and hence – as periodic waves – through the medium, the so-called luminiferous ether. And yet accepted theories in science now tell us that there is no such thing as an elastic material medium filling the whole of space – no such thing as the luminiferous ether as envisaged by Fresnel. Long before the end of the nineteenth-century, scientists had accepted instead – on the basis of its even greater predictive success – a theory that entails that Fresnel’s elastic ether is non-existent. According to Maxwell’s theory, visible light consists instead of electromagnetic radiation from a very narrow range of frequencies. Although Maxwell himself continued to believe that his electromagnetic field was ‘underpinned’ by a Fresnel-style ether, the quickly-accepted ‘mature’ version of Maxwell’s theory holds that the field is *sui generis*: it is, according to the theory, just a primitive fact about the universe that at each point in *empty* space there are, at any instant, electric and magnetic force vectors of various strengths, and light waves consist, not of oscillations of material particles about their positions of equilibrium (there are no such particles, there is no elastic solid luminiferous ether), but rather of periodic changes in the electric and magnetic field strengths (which are, to repeat, *sui generis* or primitive).

There is, of course, a range of other cases where a successful – *predictively* successful – theory has been replaced by another that is inconsistent with its fundamental theoretical claims. The replacement of Newtonian theory by Relativity Theory and of Classical Physics by Quantum Mechanics, and indeed the later replacement of Maxwell’s theory by the Photon Theory and theories of the Quantum Field are prime examples.

I see no added-value in trying to turn this historical evidence of theory-change into a ‘meta-inductive argument’ to the conclusion that our currently best theories will (probably? almost certainly?) be replaced by still better theories that contradict basic claims of those currently best theories. (The so-called Pessimistic Meta-Induction (PMI) (See Laudan [1981])) Instead I think that these historical cases should be viewed simply and directly as posing a *challenge* to the

realist who wants to claim that there is reason to hold our currently best *and fundamental* theories to be at least approximately true.

Challenge: Give me a reason to think that we are in a different situation from Fresnel, that is, give me a reason for effectively discounting the possibility that our current fundamental theories might be replaced by future theories which entail that those current theories are false (though of course still highly empirically adequate in some domains). Or, if you can't give me any such reason, explain to me how you can justify a realist attitude towards the theories accepted by current science.

Before considering ways in which that challenge might be addressed, it is important to get clearer about the extent of the challenge, about the extent of theory-change that can reasonably be claimed to have occurred in science. Even if we accept that the switch from Fresnel to Maxwell, or from Newton to Einstein, was a 'scientific revolution', no one should believe that, in a revolution, 'everything changes' (though many have taken it that this is what Kuhn was teaching us). What changes in the switch from Fresnel to Maxwell is our *fundamental* view of what light is, but theories lower down the hierarchy, the theories of optical diffraction and of polarisation for instance, *so long as they are structuralised*, remain unchanged. Similarly, although our fundamental view of matter has changed, and changed several times, since, say the early twentieth-century, science still holds, as it did then, that atoms consists of electrons, protons and neutrons. Of course, what electrons, neutrons and protons are and what their properties are, and what 'consisting of' involves are all now taken to be very different than was earlier supposed. So 'atoms consist of electrons, protons and neutrons' means something very different, if you like, now than it did then, but science still asserts the truth (and not just approximate truth) of that sentence (because that claim is implicitly taken in structuralised form). In fact, any number of 'lower-level theories' – 'the heart pumps blood around the body', 'water molecules consist of two molecules of hydrogen and one of oxygen', 'light has a finite velocity' – have been, so far as their assertability as true goes, entirely unaffected by theory-change. Again, to see Harvey's theory, for example, as retained we need to abstract from science's view of what a heart fundamentally consists of – currently an interacting system of 'ripples' in the quantum field – which has been, and is, subject to change. But the only reasonable view seems to be (again because we implicitly structuralise) that those changes do not at all affect the truth of claims like 'the heart pumps blood around the body'. Similarly, science's view of what light fundamentally is has undeniably undergone major change over history, but, the lower level theory that 'light [whatever it may fundamentally be] has a finite velocity' has been, ever since it was established in the seventeenth-century, a stable part of science (this will prove important again when we come to consider "entity realism" in section three).

Some philosophers, usually citing Larry Laudan [1981], take the premise of the PMI to be something like the claim that 'most' ('nearly all?') theories that were successful at earlier stages of science have subsequently been rejected as false. (A recent example is Fahrback [2011].) I already indicated that I see no advantage in trading in meta-inductive arguments, but even aside from that, and aside from the fact that theories are logically interrelated and so it is not easy to individuate them to perform the counts, the claim about 'most' once successful theories being subsequently rejected is – if not just fundamental theories but also those further down the hierarchies are included – patently unsustainable. The threat to realism comes, not from 'all' theories, but just from those theories that were, at least for the time being, fundamental. Fresnel's theory of optical diffraction, meaning just his account of how light, whatever it may fundamentally be, behaves when passed through various opaque screens involving slits, lives on – at least as applied to the sorts of cases for which it was developed (so, assuming for example,

that there is no significant electric field affecting the light). His theory of what light fundamentally is – a disturbance in an all-pervading elastic solid medium – has been rejected. It is the historical fact that such fundamental theories have, despite their initial success, subsequently been overthrown, to which the realist must respond.

So, having clarified the *challenge*, let's turn back to the issue of how the realist might meet it. Since the *challenge* involves an exclusive disjunction, it follows that there are two distinct (though together exhaustive) types of response that the realist might potentially make (though variation is of course possible - and indeed actual - within the types):

1. The first type of response to the *challenge* is to attempt to give a reason for holding that we are not in the same situation as Fresnel. This requires a defence of the view that our current theories are better than previously successful ones in some way that contrasts with Fresnel's theory, and that makes the eventual replacement of our current theories *at least* highly improbable, perhaps so improbable as to justify ruling out the possibility of replacement altogether.

2. The second type of response to the *challenge* is to concede that the status of the theories that are fundamental in current science is not qualitatively different from the status of the theory of light in the mid-nineteenth-century, but to deny that this creates any problem for a sensible realist. Those defending this second type of response must argue that accepting that our current theories may well suffer the same fate as did Fresnel's does not rule out reasonably holding that those current theories have latched on, at least in some approximate way, to the nature of reality. And this in turn requires an argument as to why some sort of realist attitude remains tenable towards Fresnel's and other predictively successful theories despite the fact that they are now rejected.

The first type of response has proved popular – to me at least, surprisingly so. An early first-type responder was Peter Lipton [(2000)]. The initial move behind this response is uncontroversial: *of course*, our current theories are better than the ones they replaced, otherwise they would not have replaced them. Lipton pointed out that we are justified in relying on inductive inference (if at all, I would add) only if we have no reason to think that the underlying system about which we are making the inference is changing over time. Otherwise we risk acting like Russell's about-to-be-headless chickens. And, as Lipton insists – obviously correctly – there *has* been change since Fresnel's time: the theories accepted in science now are different from, and of course better than, those accepted in Fresnel's time. But this at most represents an argument against thinking of the history of theory-change in science as the basis for a 'meta-induction', which, I already suggested, would be unhelpful. Even if we agreed with Lipton that the fact that our current theories are better than their predecessors makes inductive inference to the likelihood of eventual change of our current theories inappropriate, that would by no means meet the *challenge* posed by theory-change. A successful response to the challenge surely requires something more, namely, a demonstration that there is some *qualitative*, rather than merely quantitative, difference between the theories that are currently accepted in physics and their successful but now rejected predecessors. However, so long as we restrict ourselves to genuinely predictively successful theories (as Laudan failed to do when constructing his notoriously inflated list of 'successful' theories now regarded as false!), then it is difficult to see any contender for such a qualitative difference. Current theories of light based on Quantum Field Theory and the photon are better than Maxwell's electromagnetic theory, for example, but not in any qualitatively different way than Maxwell's theory was better than Fresnel's: current theories 'only' correctly predict all the phenomena that Maxwell's theory predicted plus some more, exactly as Maxwell's theory had predicted all the phenomena that Fresnel's theory successfully

predicted plus some more. Arguing that previous rejections of successful (fundamental) theories create no *prima facie* reason to take seriously the prospect of the future replacement of current theories because our current theories are better than their predecessors is, it seems to me, exactly on a par with arguing that, even though previous world records have been broken, there is no serious chance that Usain Bolt's record for the 100m will ever be broken because it is significantly better (i.e. shorter) than previous records! Remember that scientists in the nineteenth-century really did believe that Newton had once and for all discovered the truth about the universe.

The most recent philosopher of science to endorse the first-type of response is Ludwig Fahrbach (see his [2011] and also his [2017].) Fahrbach identifies an 'exponential growth in science' reflected in the (alleged) fact that 'at least 80% of all scientific work ever done has been done since 1950' ([2011], p. 139). And he goes on to claim that that 'practically all of our most successful theories were entirely stable during that period of time [i.e. since 1950]' (*ibid.*) He concludes that theory-change is at least very largely an historical phenomenon which therefore presents no reason to doubt the realist credentials of *current* (post-1950) science.

There are a number of ways in which Fahrbach's analysis is misguided. He bases his claim about the 'exponential growth' of science in (large) part on the growth of the number of journals, and journal articles devoted to scientific issues over time. These numbers surely (a) reflect social and economic factors, rather than genuinely scientific ones and (b) mean that most of the pieces of 'scientific work' counted are very much in the realm of Kuhnian normal science. This is reflected in the list of theories supplied by Fahrbach (*ibid.*, p. 152) to substantiate his claim that 'practically all of our most successful theories were entirely stable' since 1950. The list includes 'The Periodic Table' (hardly a theory, but never mind) and the theory that 'The brain is a net of neurons'. As explained earlier, this sort of lower-level theory is irrelevant to the real issue facing realism which concerns changes in (pro tem) fundamental theory. Fahrbach does at one point address the hackneyed case of changes in theories of light, and therefore the real issue. He writes (*ibid.*, p. 149): 'As regards the example of theories of light, all changes of those theories occurred before the 1930s, whereas 80% of all scientific work ever done has been done since 1950 [...] Thus it seems that the set of examples offered by anti-realists is not representative and cannot be used to support' any challenge to which the scientific realist needs to respond. But the idea that changes in our fundamental view of the constitution of light stopped before 1930 is (very) false. It is not at all clear why Fahrbach should have chosen 1930 as the critical year, but developments in basic light theory did not, of course, stop with Einstein's postulation of discrete photons in 1905. Our whole view of, if you like, what photons are, has been radically changed first by developments of Quantum Electrodynamics and then of Quantum Field Theory. Any serious history of science since 1930 would replace Fahrbach's stable view with a ferment of ever-changing (and still very far from settled) fundamental theories of light.¹

It seems to me then that only the second type of response to the *challenge* posed by theory-change remains open to the realist. Any such response involves conceding that it is indeed perfectly possible that our current fundamental theories, for all their success, will be replaced by so-far undreamt-of, still more successful theories; but then proceeding to argue that this concession is *not* incompatible with reasonable belief in the theories of current science as, in some sense, approximately true. This in turn surely requires a demonstration that the relationship between successive successful fundamental theories in the historical record is *not*

¹ Even if things in light theory *had* been quiet since 1930, weren't things quiet in (fundamental) mechanics from 1666 to 1905?

one of 'radical rejection', as supposed by those who see theory-change as directly at odds with any version of realism; but instead that the record is one in which the earlier theory (if, remember, predictively successful) continues to be somehow reflected in the later one (beyond, of course, simply having its true empirical consequences retained). This would mean that that earlier theory is, in some sense, 'partially or approximately retained' within the later one.

The most natural version of this second type of response involves interpreting the claim that the earlier theory is 'partially retained' in its successor as meaning exactly that (substantial) parts of the earlier theory are literally retained within its successor. This is Stathis Psillos's *divide et impera* manoeuvre ([1991], pp. 103–109): in the hackneyed example, the claim would be that some parts of Fresnel's classical elastic solid ether wave theory are retained within Maxwell's 'mature' theory of the electromagnetic field – the very parts that were responsible for the success of Fresnel's theory. Indeed, it seems so natural to interpret 'partial retention' in this way that it is widely, if only implicitly, assumed that *any* type-two response to the challenge is bound to be some version of this *divide et impera* or *selective realism* view. However, this is not correct, as we shall now see.

What other account of 'partial retention' is available? My claim that *divide et impera* is not the only strategy available to those making the second type of response to the challenge of theory-change requires that some more general notion of 'partial retention' be available. I suggest that, in fact, the basic requirement for an earlier theory to 'continue to be reflected in' or to 'partially or approximately retained within' its successor theory is simply that the predictive success of the earlier theory be explicable from the vantage point of the later theory – explicable as in no way a 'miracle' but rather as what was to be expected given the claims made by the earlier theory. So, the basic desideratum for a successful type-two response is that it should provide the template for such explanations. After all, the motivation behind any realist view is the 'no miracles' intuition and exactly such an explanation is what is required to defend that intuition against the challenge posed by theory-change. In the case of our hackneyed example, the basic requirement for a successful type-two response is, therefore, on this account, a 'non-miraculous' explanation from the vantage point of Maxwell's theory for the striking predictive success of Fresnel's theory. One way, but *only* one way, of providing such an explanation would be *via* the claim that substantial parts of Fresnel's theory were in fact carried over into Maxwell's, and that the parts that were carried over were the ones responsible for the predictive successes: if that were true, then no wonder that Fresnel's empirical predictions turned out to be correct, his theory was – in substantial part, though not of course wholly – correct. (Of course, correct when judged from the vantage point of Maxwell's theory). And the predictions could be derived from the 'correct' parts alone. But, to repeat, I shall take it that *any* such explanation is enough, if successful, to defend the no miracles intuition and hence to provide a defensible version of realism that overcomes the challenge from theory-change. And as we shall see, the type-two response that I endorse – structural realism – gives a different type of explanation of the success of originally successful but now rejected theory and is therefore, contrary to a widely held view, not accurately regarded as a version of selective realism.

Where there is such an explanation for the predictive success of the earlier theory from the vantage point of the later, it seems reasonable to take it that the earlier theory continues to look 'approximately true' from the vantage point of the later theory. Maybe a stronger, more formal, notion can be developed of what it means for one theory to appear to be 'approximately true' in the light of another, but I doubt it, and I will certainly mean no more than this when I use the term 'approximately true' in this paper. 'Fresnel's theory continues to look approximately true from the point of view of Maxwell's theory' is then, for me, just shorthand for 'the success of

Fresnel's theory is in some way explicable from the vantage point of Maxwell's' (where the 'explanation' that it simply happened to have some of the same correct empirical consequences, does not count). If it can be shown that historically all replacements of predictively successful theories by still better ones have been cases of replacement in which the earlier theory continued to look approximately true in this very weak sense then, if meta-inductive arguments are to be traded in at all (which I advise against), the history of theory-change in science will form the basis for an *optimistic* meta-induction to the conclusion that, although our current theories may well eventually be rejected in favour of still better ones (and those in their turn rejected in favour of even better theories), our current theories will always look, from any of those later vantage points, to an important degree correct – correct in their trans-phenomenal and just their phenomenal claims. And I take this to be the central claim that any version of scientific realism must assert and defend.

So how can a realist substantiate the claim that earlier theories, if predictively successful, continue to look 'approximately true' from the vantage point of their successors in the sense just indicated? Here, using the motivating example we have discussed and in the words of its great originator, Henri Poincaré, writing in 1905, is what I hold to be the uniquely successful way – so-called (epistemic) structural realism (hereafter 'SR'):

The ephemeral nature of scientific theories takes by surprise the man of the world. Their brief period of prosperity ended, he sees them abandoned one after the other; he sees ruins piled upon ruins; he predicts that the theories in fashion today will in a short time succumb in their turn, and he concludes that they are absolutely in vain. This is what he calls the *bankruptcy of science*.

His scepticism is superficial; he does not take account of the object of scientific theories and the part they play, or he would understand that the ruins may still be good for something. No theory seemed established on firmer ground than Fresnel's, which attributed light to the movements of the ether. Then if Maxwell's theory is preferred today, does it mean that Fresnel's work was in vain? No, for Fresnel's object was not to know whether there really is an ether, if it is or is not formed of atoms, if these atoms really move in this way or that; his object was to predict optical phenomena.

This Fresnel's theory enables us to do today as well as it did before Maxwell's time. The differential equations are always true, they may always be integrated by the same methods and the results of this integration still preserve their value. It cannot be said that this is reducing physical theories to practical recipes; these equations express relations, and if the equations remain true, it is because the relations preserve their reality. They teach us now, as they did then, that there is such and such a relation between this thing and that; only the something which we then called *motion*, we now call *electric current*.

But these are merely names of the images which we substituted for the real objects which Nature will hide for ever from our eyes. The true relations between these real objects are the only reality we can attain... ([1952], pp. 160–161, original emphasis).

SR will, of course, be centre-stage in what follows, but, for now, notice three things about this passage: (1) Poincaré, as the first paragraph shows, was completely aware of the challenge posed for realism by the history of theory-change (many decades before Laudan) and, as the second paragraph shows, believed he could meet that challenge; (2) although Poincaré is often

considered to be an anti-realist instrumentalist, the third paragraph, in which he explicitly rejects the claim that his view would 'reduc[e] physical theories to practical recipes', clearly contradicts this; and (3) the positive view, SR, is contained in the third and fourth paragraphs – notice in particular that the final (fourth) paragraph makes it clear that the view being defended is an epistemic rather than an ontological ('ontic') claim: the 'relations' do not (at least do not necessarily) exhaust reality, but they are 'the only reality we can attain' that is, they are all that is epistemically accessible. And we can be realist about 'relations' (or structure) despite theory-change. Moreover, since the relations are the 'only reality we can attain', the effective content of any theory is just its mathematical structure (the way it structures the phenomena) – the rest is just 'names of images...[that] Nature will hide forever from our eyes'. This means that the 'ruins' seen by the (remarkably sophisticated) 'man of the world' are in the end a mirage, the 'ruins' consist of metaphysical/heuristic images with no real cognitive content. This is why SR is not a version of selective realism. More of this soon.

This current paper constitutes a further defence of SR – first by comparing it with important rivals. In the next section of the paper I will consider what I already described as the most natural way of construing claims like 'The success of Fresnel's theory can be explained from the vantage point of Maxwell's theory.' This is to interpret the claim in the 'selective realist' way espoused by both Philip Kitcher ([1993] and [2001]) and Stathis Psillos *op. cit.*): parts of Fresnel's theory were carried over into Maxwell's; and those parts were anyway precisely those responsible for the predictive success of Fresnel's theory, while, on the contrary, the parts of Fresnel's theory that were rejected as a consequence of the shift to Maxwell played no role in the predictive success of Fresnel's theory. I argue that their selective realist account fails and at the same time explain further why SR is not the version of selective realism that many believe it to be.

Other realisms aside from selective and structural are, of course, also on the market. Perhaps the *prima facie* strongest and certainly the most eminently supported is the 'entity realism' of Ian Hacking and Nancy Cartwright. In the third section of this paper, I analyse entity realism. I argue that, although not presented as a (type-two) response to the challenge of theory-change, it is readily revealed on analysis to constitute such a response and an inadequate one at that.

Suppose I am right that neither selective nor entity realism is a viable position, then a tempting suggestion – in view of what many have seen as fatal problems for SR – is that *no* realist position is viable: that there is in fact no viable realist response to the challenge posed by the history of theory-change in science. In the final section (section 4) I examine what I think are the two main issues that many have seen as posing insuperable objections to SR and argue that they can in fact readily be countered. If I am right, then SR is viable; and indeed, in view of the untenability of either selective or entity realism, SR would seem likely to be the *only* viable realist game in town.

2 Selective Realism and its Problems

There is, as already noted and even ahead of any consideration of theory-change, no realist warrant for *provably redundant* assumptions such as Newton's hypothesis about the centre of mass of the solar system being at rest in absolute space. (And indeed, adding to what Kitcher and Psillos claim, there is clearly no realist warrant for non-redundant but blatantly *ad hoc* assumptions such as Copernicus's assumption of a 'Third Motion' of the earth.) So, any sensible realist is automatically selective as regards obviously redundant or clearly *ad hoc* (non-

independently-testable) parts of theories. However, the sort of theoretical claim whose later rejection poses the challenge for the scientific realist is not like Newton's (or Copernicus's). Fresnel's claim that the medium that carries light waves is a highly attenuated elastic solid, for example, is not obviously redundant in the way that Newton's 'hypothesis' is – there is certainly no proof that altering Fresnel's assumption at will yields an overall theory which still has all the same empirical consequences as the original. (Of course, this means only that there is no such proof that does not adopt the vantage point of later theories: needless to say, Maxwell's theory (or what I called the 'mature version' of it) shows that Fresnel's notion of the ether is 'redundant' in the sense that you can get exactly the same effects (and more) from a *sui generis* electromagnetic field.) And the claim that there is an ether seems, *prima facie* at least, to be a central part of Fresnel's overall theory and to figure in its success.

Yet both Kitcher and Psillos base their version of selective realism on the claim that the ether is in fact an 'idle posit', or 'non-working assumption' within Fresnel's theory; and that that is why its eventual rejection poses no problem for realism. Hence Kitcher's and Psillos's view goes significantly beyond the uncontroversial version of selectivism. Kitcher writes:

The ether [in Fresnel's theory] is a prime example of a presuppositional posit, rarely [sic] employed in explanation or prediction, never [sic] subjected to empirical measurement [...] yet seemingly required to exist if the claims about electromagnetic and light waves were to be true ([1993], p. 149).

This already sounds rather unclear. Since 'subjecting a supposed theoretical entity to empirical measurement' is just another way of saying that the claim that the entity exists made predictions that have been successfully tested, we are entitled ask 'which is it: did the "ether" do real predictive scientific work rarely or never?' If rarely, then presumably it was somewhat involved in the success of the theory and so not entirely redundant. But there is no need to pursue this issue further since neither 'rarely' nor 'never' is correct, as we shall almost immediately see.

So, the 'ether' is alleged to be redundant on Kitcher's account, while, by sharp contrast, the notion of a 'light wave' is a 'working posit' – one that was directly involved in the predictive success of Fresnel's theory. And, exactly as the selective realist, would expect, Fresnel's notion of a 'light wave' is carried over into, and indeed continues to be judged as referential in the light of, Maxwell's replacing theory – or so claims Kitcher (and also Psillos).

Neither the claim about the ether nor the claim about light waves is sustainable. Let's first examine the claim about the redundancy of the ether. The notion of 'redundancy' at issue here is a difficult one in general, as both proponents of this view at least partially acknowledge. Craig's lemma (explicitly referred to by Psillos (*op. cit.*, pp. 20–4)) suggests that, unless we are careful, *every* theoretical claim can be declared 'redundant' to the empirical success of the theory.

Moreover, earlier work in philosophy of science shows that natural attempts to formalise the notion of redundancy make it a worryingly axiomatisation-dependent notion. Invoking the notion of redundancy seemed, for example, the obvious way to resolve the 'tacking paradox' that afflicts some accounts of theory-confirmation. The General Theory of Relativity (GTR) conjoined with the statement 'God exists', for example, shares all the empirical consequences of GTR alone (in the presence of appropriate auxiliaries, of course). But no one would claim that the conjunction 'GTR & God exists' is confirmed by the correctness of those empirical consequences – at least not in the sense of making its claim about the existence of God better empirically founded. The natural explanation of that judgment about confirmation seems to be that, since 'God exists' is *redundant* to the empirical success of the conjunctive theory – that is, since we can

derive all the empirical successes just from GTR without invoking the claim that God exists - confirmation does not spread to the latter claim. Confirmation, it seems natural to say, spreads only to theoretical assumptions that are necessarily involved in the derivation of empirical consequences that turn out to be correct. And it is indeed true that, for 'natural axiomatisations' of the theory T' (GTR & God exists), the claim that God exists is redundant in this sense. Suppose, to illustrate, that we could axiomatize GTR (T) using, say, five axioms A_1, \dots, A_5 , then the natural axiomatization of T' is just A_1, \dots, A_5, G where G is the statement 'God exists'. Relative to those axiomatizations of the two theories (T and T'), it is indeed true that every empirically testable consequence derivable from T' is already derivable from T , and so the assumption G seems redundant to the empirical success of T' - just as we would intuitively expect. But T' could, of course, be re-axiomatized, as, for example, $G \rightarrow A_1, \dots, G \rightarrow A_5, G$. This re-axiomatization is obviously logically equivalent to the original axiomatisation A_1, \dots, A_5, G - they are both axiomatisations of T' . But within that re-axiomatisation, G , far from being redundant to the empirical success of theory, is crucial: every single derivation of an empirical consequence requires the axiom G .

So, redundancy is by no means a straightforward notion. But even laying aside such general logical problems, the specific assertion that the ether was a redundant part of Fresnel's theory is difficult to sustain. The fact that light waves must pass through a medium *of some sort* is indeed, far from being redundant, central to the theory: the light vibrations are created in luminous sources such as the sun, a finite time later they are received by some light sensitive apparatus: the human eye, say - the energy has to be stored in, be carried by, *some sort of medium* in the interim. Moreover, Fresnel explicitly talked of 'particles' of the ether and got mathematical mileage out of the assumption that, when taken away from their position of equilibrium as a result of the passage of light, they are subject to an elastic restoring force, obeying Hooke's law. Fresnel's assertion that the ether exists is significantly dis-analogous to Newton's 'hypothesis' about the centre of mass of the universe being at rest.

But things become even more problematic for this selective realist view when we turn to Kitcher's claim that Fresnel's notion of a 'light wave' is, in direct contrast to his notion of the ether, a working posit and that correspondingly, just as the selective realist would expect, Fresnel's notion of a light wave was carried over into Maxwell's superseding theory (and later ones).

But what *was* Fresnel's notion? That is, what did Fresnel hold a light wave to be. Well, if you were to imagine freezing the light wave at some instant, it would consist, according to the theory endorsed by Fresnel, of the trace of material particles of the ether at different distances from their positions of equilibrium, and hence with different kinetic energies and, correspondingly, subject to different values of the elastic restoring force. In order to maintain the view that *Fresnel's* notion of a light wave (as opposed to some other more general notion) was retained within electromagnetic theory, you already have to structuralise that notion: to abstract from what it was that Fresnel himself thought was waving and consider the waves only as periodic motions *of some quantity or other*, motions that satisfy certain mathematical equations.

Kitcher sometimes seems fully aware of this. He writes, for example:

Fresnel's success is based on his insights into the propagation of transverse waves, and his faulty beliefs about the *constitution* of those waves are irrelevant to his analysis. This, of course, is why his work has been so successfully encapsulated in contemporary physics. ([1993], p. 147)

Structural realism has no truck with the claim, which as we just saw is historically false, that the constitution of the light waves, as waves in an elastic solid ether, was 'irrelevant to [Fresnel's] analysis'. But SR does assert that 'only' a structural facsimile of Fresnel's ideas about light waves was retained within later theories of light. And Kitcher seems here to be agreeing with this. So, Fresnel's notion of 'light wave' needs to be structuralised if it is to be seen as retained in later theories of light. However, Fresnel's notion of the ether can just as easily be structuralised – as being the medium, *whatever it turns out to be*, that carries the waves whose description satisfies various mathematical constraints. And the ether, as thus structuralised, was equally retained in later theories. The terms 'ether' and 'light wave' in Fresnel's theory stand on a par – which is just as it should be (they are inseparable from one another in Fresnel's thought) and just as SR suggests. Selective realism fails to draw the distinction that it needs to between the ether and light waves. And it inevitably fails: there is no *selective* rejection/retention involved in the shift from Fresnel to Maxwell, instead a structural facsimile of the whole of Fresnel's theory is retained in Maxwell.

To repeat: what is retained of Fresnel's theory within electromagnetic theory is not a *part* of the theory but rather a structural facsimile of the whole theory. Indeed, it could not, on reflection, be otherwise: Fresnel's notion of the ether is intimately connected with his notion of a light wave, which is indeed built upon it.

At least then for the paradigm case, the selective view does not work – the correct account of the relationship between Fresnel's and Maxwell's theories is the one in line with SR. The attempt to argue the selective realist case in fact leads on analysis to SR. Yet Kitcher not only holds that selective realism is different from SR, he sees it as logically stronger than SR: he argues that SR is as unnecessarily weak – it makes, as he sees it, more concessions to the facts about theory-change than is necessary. Why? Kitcher writes:

Realists [should] ponder why Fresnel's theory was so successful, given that his beliefs in the ether and in the identification of light with waves in the ether (the jostling of ether particles) were wrong. An obvious place from which to start is the contemporary optics textbook, where Fresnel's mathematics of wave propagation is reproduced. So one might propose that Fresnel was right about the *structure* of the mechanisms that produce the phenomena with which his theory dealt so successfully – diffraction, interference, and so forth – even though he was wrong about the *substance*, linking his mathematical equations (the correct structure descriptions) to a faulty ontology (by taking the symbols to stand for conformations of the ether). From this diagnosis of the Fresnel example, one might then proceed to recast realism to make it more modest: success turns out to be a reliable indicator of underlying structure, and realists are only entitled to assert that successful theories are correct in delineating unobservable structures. ([2001], p. 169, original emphasis)

But then, having expressed SR so clearly, Kitcher goes on to reject it as unnecessarily weak: 'But that, I believe, is too great a concession [to the facts about theory-change] ... there is no principled reason to think that the reliability of our [scientific] inferences is restricted to structural matters.' (*ibid.*)

Kitcher in fact sees it as a central feature of his selective realism that it allows that, at least some of Fresnel's theoretical terms, including that of 'light wave', have real world reference. Whereas, as he correctly presumes, SR makes no claims about 'continuity of reference' for basic theoretical terms across theory-change. So, it is the alleged underwriting of continuity of reference for at least some of Fresnel's theoretical terms that marks selective realism as a

logically stronger view than SR. We have already seen that selective realism evaporates on analysis, but the issue about reference is independently interesting.

The claim that the theoretical terms in our successful theories have real world reference has, of course, standardly been taken to be central to scientific realism: the realist must first (and foremost) assert that the terms in the successful theories in current science refer to 'real world' entities (or rather that we have good grounds for believing that they do). Given that the facts about theory-change make the claim that our current theories are literally true so implausible, the realist is allowed the leeway to claim only that what those current theories say about those real world entities is approximately true. The realist must claim that electrons, for example, are 'real', even though what our current theories say about electrons is only approximately true. So Laudan, for instance, was clear that the claim that at least 'central' terms in current theories refer is essential for the realist: 'I take it that *a realist would never want to say that a theory was approximately true if its central terms failed to refer*'. ([1981], p. 33, original emphasis) But Laudan's attack on realism was, of course, explicitly directed at the sustainability of the claim that central terms in previously accepted theories refer. So if indeed the realist is committed to this claim about reference then she faces a particularly challenging version of the challenge from theory-change.

Challenge 2: Assuming that our current theories are not sacrosanct (that is, assuming a type-two response to the original challenge), the 'referential realist' needs an account that allows theoretical terms to be referential even though the theories in which those terms occur are not strictly true; and, in order for such an account to be plausible, it ought presumably to sanction terms in previously accepted but now superseded theories as referential (not, notice, 'approximately referential' whatever that might mean).

Kitcher, as noted, takes on this second challenge – albeit only with regard to theoretical terms involved in 'non-idle', 'working' parts of successful theories.

With regard to the hackneyed example, Kitcher claims concerning reference are, remember, that:

1. The term 'ether' as used in Fresnel's theory is non-referential but that's alright for the selective realist since the ether was an 'idle posit'.
2. The term 'light wave', on the contrary, features in the working part of Fresnel's theory but that's alright too for the selective realist, because 'light wave' is a referring term.

Kitcher writes:

Fresnel's dominant intention is to talk about light, and the wavelike propagation of light, however that is constituted. He has, of course, a false belief about the medium of propagation. But, since his primary aim is to discuss light and its wavelike qualities, his tokens of 'light wave' in the solutions of the diffraction and interference problems genuinely refer to electromagnetic waves of high frequency. ([1993], p. 146)

Kitcher makes this claim despite the fact that Fresnel of course went to his death bed without having the slightest idea of what an electromagnetic wave might be.

There has been a good deal of criticism of Kitcher's account of reference. Reference is surely supposed to be a relation between a linguistic item and the world – bringing in speakers' intentions means that you are analysing some other notion, not reference as classically conceived. But the two most important problems for the account for present purposes can both

be seen without challenging Kitcher's account of reference in general or the role of speakers' 'intentions' within it.

First, that account cannot do what Kitcher thinks it needs to do – namely differentiate between the allegedly referential term 'light wave' and the allegedly non-referential term 'ether'. Notice again how Kitcher argues for the term 'light wave' being referential in the passage just quoted from his [1993]. The problem (as we would expect from what I noted earlier) is that an entirely analogous argument can readily be constructed for the referential nature of the term 'ether' (an argument which I place between asterisks for ease of later reference(!)):

* Fresnel's dominant intention is to talk about the medium that carries the waves of light, however that is constituted. He has of course a false belief about the precise constitution of that medium, but since his primary aim is to discuss light and its wavelike transmission, his tokens of 'ether' in his work on diffraction and interference problems genuinely refer to the electromagnetic field.

Kitcher explicitly claims, as we have noted several times, that the term 'ether' within Fresnel's theory is non-referential. So, he must – implicitly – hold that there is a difference between the claim about light waves quoted above and the analogous claim about the ether that I just constructed and placed within asterisks: because he defends the first and rejects the second. But I can't myself see what that difference could possibly be: the terms 'ether' and 'light wave' in Fresnel stand on a par – they are, as noted, inseparable from one another in Fresnel's thought. Parity between the two is exactly what SR entails as we saw already: there is no *selective* rejection/retention involved in the shift from Fresnel to Maxwell; instead a structural facsimile of the *whole* of Fresnel's theory is retained as part of Maxwell's.

It seems that that only two positions are tenable (a) *neither* 'light wave' *nor* 'ether' (as understood by Fresnel) refers (which is the SR position) *or* (b) if Kitcher wants to continue to insist that we can have a stronger version of realism than SR involving reference, he will have to argue that *both* 'light wave' *and* 'ether' (in Fresnel's usage) refer. In fact, just as suggested in my asterisked passage, Kitcher will have to claim that the term 'ether' in Fresnel's theory does indeed refer – it refers to 'the electromagnetic field'.

This suggestion about the reference of the term 'ether' in Fresnel has proved independently attractive to some philosophers. It was, for example, defended in a much-cited brief article by Hardin and Rosenberg ([1982]). The suggestion seems to me, however, entirely untenable. First, it seems odd to have Fresnel refer in 1818 to something of which he had not, and could not have, the slightest conception – since the idea did not arise until 50 years later. And it seems *distinctly* odd to have him committed to statements about the entity that he allegedly referred to that are quite radically false, such as that the electromagnetic field (in the mature Maxwellian sense) consists of an elastic solid. (Putnam, of course, famously advocated the exercise of a 'principle of benefit of doubt' ([1975b], p. 274) in judging whether the terms of earlier, now superseded theories refer: but an immaterial electro-magnetic field seems not to be just 'a little bit different' from a material elastic solid!)

Much more importantly, those who would have Fresnel refer to the electromagnetic field when talking about the ether have not really faced up to the full challenge of the history of theory-change in science and in particular have not reflected properly on the fact that theory-change in optics did not come to a halt with the shift from Fresnel to Maxwell. If the suggestion means, as Hardin and Rosenberg assume, that Fresnel's term 'ether' refers to the electromagnetic field *as characterised in the mature version of Maxwell's theory* then, Fresnel's

term remains non-referential by current lights. Science no longer accepts the claim that the electromagnetic field in Maxwell's exact sense exists. We no longer think of the 'electromagnetic field' as precisely characterised in the mature version of Maxwell's theory as a referential term – any more than the term 'ether' as characterised in Fresnel's theory. Current theories of light involve photons in a strange quantum field obeying probabilistic laws, quite different from anything in Maxwell (though of course connected to many phenomena in structurally analogous ways). It seems that what Hardin and Rosenberg (and second-horn-of-the-dilemma Kitcher) would have to say is that in using the term 'ether' Fresnel was 'really' referring to the probabilistic quantum field. But even that would only be a judgement about what Fresnel was 'really talking' about *from the vantage point of our current theories*. But this whole debate, at any rate in terms of type-two responses to the (initial) challenge of theory-change, is, remember, predicated on the assumption that our current theories might well themselves be replaced by theories currently beyond our ken. So, Quantum Field Theory may well itself later be superseded. Hence, in the end Kitcher and Hardin and Rosenberg would have to admit that we don't know what Fresnel's term 'ether' referred to! Except that it is whatever it is (if it is, indeed, one unified entity) that underlies optical and related phenomena. And then they would be exactly correct because they would have adopted SR.

The attempt to invoke reference, in its classical semantic sense, seems to be based on a yearning to embrace an assumption that everyone explicitly rejects when articulated but nonetheless often find themselves implicitly adopting: the assumption that we have some direct non-theory-mediated access to reality that we can compare available theories to. Stripped of this clearly untenable (though constantly beguiling) assumption, reference for theoretical terms is nothing special. Statements such as, for example, 'there are electrons' or 'there is an electromagnetic field' are just individual, existentially generalised consequences of accepted theories. Once we understand the term 'electron' to mean '(potential) entity as, if it exists, characterised by our current best theory of electrons', then those aware of theory-change in science will immediately recognise that 'there are electrons' (in that sense) stands or falls with that current theory; and that, since that theory may well eventually be superseded, claims of reference are on a par with, are just as tentative and revisable as, claims about truth. Putnam really did start the debate off down a blind alley.

'Reference' when understood in the only way appropriate for theoretical terms is just as tentative and approximative as truth. The only defensible version of realism seems to be 'realism without reference' – i.e. SR (Papineau and Cruse [2002]). What Kitcher, explicitly, and Hardin and Rosenberg are in fact best construed as providing is not an account of reference in the classical sense – as an (alleged) correspondence between a linguistic item and reality, but instead a suggestion for a translation between terms in two different theories. 'What notion in Maxwell's theory corresponds most closely to Fresnel's notion of a light wave as a disturbance in a mechanical medium?' Answer: an electromagnetic wave from within a particular band of frequencies. 'What notion in Maxwell's theory corresponds most closely to Fresnel's notion of the elastic solid ether?' Answer: the electromagnetic field.

Or, if we want to honour the original intention of the classical notion of reference as involving the real world, then Kitcher's account, abstracting from the 'intentions' and 'particular tokens' idea, can be thought of as answering the following conditional question 'What would be the real-world entity that comes closest to satisfying Fresnel's notion of a wave in a mechanical medium, on the assumption (which anyone sensible would regard as counterfactual) that the superseding theory (or the current theory, whichever later vantage point we want to take) actually "carves nature at the joints" and hence tells us what the real world entities are?' Similarly, Hardin and

Rosenberg are really saying ‘What would be the real-world entity that comes closest to satisfying Fresnel’s notion of an elastic solid ether, on the assumption (which we already know to be incorrect!) that Maxwell’s theory is true and so its terms refer?’ These are all (inevitably) questions about the logical relationships between theories, not questions about the relationships between any theory and ‘the real world’. We can’t secure reference for our theoretical terms and can endorse a version of scientific realism without needing to secure reference. The claim that Fresnel’s notion of ‘light waves’ refer is just an (unnecessary) meta-level rendition of ‘there are light waves (as characterised in Fresnel’s theory)’. This existential claim we know to be false. SR shows that you can accept the falsity of the claim and still be a sort of realist. Real realism, the structural kind – SR, is realism without reference. Selective realism, as advocated by Kitcher (and Psillos) cannot avail itself of a coherent notion of reference to underpin its claim to be a defensible, stronger version of realism than SR.

3 Entity Realism

Philosophers looking to purchase a version of scientific realism are of course not restricted just to the selective and structural models. Of the others, probably the most influential and certainly most eminently supported is so-called entity realism. This was initially named and defended by Ian Hacking in his 1983 book *Representing and Intervening* – which cites ideas also to be found in the work of Nancy Cartwright. A later paper of Hacking’s ([1988], p.277) provides a succinct summary of the position:

In 1983 my *Representing and Intervening* argued that nothing to do with theories (representing) helps settle the endless debates about scientific realism/anti-realism. But it urged that experimental science (intervening) strongly leads on to realism about the entities postulated by theories.

The most famous slogan associated with entity realism is of course ‘If you can spray them [electrons], then they are real’ (Hacking [1983], pp. 22–24). The idea is that, although there may be no grounds for thinking that our current full-fledged theories of electrons are true and hence no grounds for a realism about electrons as exactly construed within those high-level theories, if we can experimentally manipulate electrons by, for example, spraying them at a target then they are surely real (or, rather, then we surely have solid grounds for holding that they are real). Hacking writes:

We are completely convinced of the reality of electrons when we regularly set out to build – and often enough succeed in building – new kinds of device that use various well-understood causal properties of electrons to interfere in other more hypothetical parts of nature. (*ibid.*, p. 265, whole sentence in italics in original.)

Hacking’s book (along with related work by Nancy Cartwright and earlier work by historians such as Allan Franklin) focussed philosophical attention on the hitherto largely neglected area of experiment in science and argued that the idea that all experiments, or even all significant experiments, are tests of theories is a myth. Experiment has a life that is at least partly independent of (basic) theory. There is no doubt something to this; but I cannot see how a defensible version of realism can be based on experiment in the way that Hacking supposes.

He makes it appear as if entity realism is a different sort of animal from selective realism: rather than being a response to the *Challenge* posed by the history of theory-change in science, it (allegedly) eschews considerations based on theories ('representation') entirely and concentrates instead on what it sees as the realist-inclining power of experiment ('intervention'). I shall argue that in fact, when analysed more carefully, entity realism *is* a response to the *Challenge* – just like selective realism – and an equally unsatisfactory one (though for different reasons).

First, notice that entity realism cannot defensibly assert a biconditional claim – the claim surely cannot be that an entity is real (or better: the assertion that an entity is real is defensible) if and only if it is manipulable (really: if and only if we *believe* we can manipulate it). For example, scientists prior to the 20th century surely had very good grounds for thinking that the planets (which should surely count as theoretical entities) are real even though there was no sense in which they could manipulate them. (During the twentieth-century we learned how to 'manipulate' them at least in the sense of hitting (some of) them with projectiles.) So entity realism must be a partial view: amongst those elements that we have good reason to hold are real are those elements that we can manipulate.

Secondly, as Hacking's articulation of it makes transparent, entity realism is firmly predicated on Putnam's account of reference. The entity realist needs to be able to underwrite the claim that the 'electrons' that a twenty-first-century experimental physicist takes herself to be spraying at various targets are exactly the same things as Thomson identified as the constituents of cathode rays and Millikan attached to oil drops suspended in electric fields. Yet, as Hacking himself sharply summarizes (*ibid.*, pp. 83–84), Thomson's full theory of what electrons are, for example, is very different from Bohr's, which in turn is very different from modern conceptions involving quantum fields. Thomson held what is often called the 'pudding theory' of electrons which entailed that they are situated in atoms like raisins in a pudding. Bohr's electrons, remember, most of the time do not even have well-defined spatio-temporal positions!

While Hacking's account makes it clear that he sees weaknesses in Putnam's account of reference for theoretical terms, he clearly believes that we need to cling to it (or something like it) if we are to avoid the dreaded incommensurability. How can theory-change in science be even remotely rational if the earlier and later theories are not talking about the same things? I think that this is a mistake. We can perfectly allow that Thomson's 'electron' would, if his theory were true, be an entirely different entity than Bohr's 'electron', say and admit, indeed, that it is a purely conventional issue that we happen to continue to use the same word, without getting involved in any worrying 'incommensurability'. I will explain my reasons for rejecting Putnam's causal theory in section four along with my reasons for holding that we can do without continuity of reference across theory change while still being rationalists about theory-change and indeed while still being realists (of a structural sort). For the moment, I just note that, since entity realism requires Putnam's causal theory of reference, any objections to Putnam's theory (and there are lots, as even Hacking quietly concedes (See e.g. *ibid.* pp. 84–91)) are automatically objections to entity realism.

The most important and direct objection to entity realism, however, is the fact that theory is obviously still very much involved. The line between 'representation' and 'intervention' is not as sharp as at least the entity realist slogans suggest. To see this, we need only ask why it is that an experimental physicist believes she is 'spraying electrons' (or more generally manipulating some supposed entity) in some specific circumstances. Such an experimental physicist obviously has no knowledge of electrons 'by acquaintance' – though perhaps she should count as having acquaintance with the tracks in a cloud chamber that she takes to have been caused by electrons.

But the beliefs that she is spraying the electrons and that she is seeing the effects of the passage of an electron through a cloud chamber when she sees a particular track are, of course, beliefs based on *theories*. Those theories may not, as Hacking insists, amount to anything like everything that current frontier science attributes to electrons; she may indeed not know (or care) much about the fine detail of how electrons are regarded in, for example, Quantum Field Theory, but it will nonetheless be a substantial body of theory. So, while she may have no realist commitment to fundamental, top-level theory about electrons, the experimental physicist must have a realist attitude to those lower-level theories about 'electrons' which underwrite her assumption that, in particular circumstances, she is spraying them (whatever they really are!) at a target. 'Intervention' (or rather beliefs about when we have or have not intervened in a particular way) are (of course) somewhat dependent on representation.

The theories that the experimenter relies on are, however, lower down the hierarchy – sufficiently low to escape theory-change (as least in name). The 'electron' of current theories is, as just noted, a far cry from what J.J. Thomson supposed it to be – his 'pudding picture' of the electron has been firmly rejected – but science still includes the theory that 'electrons' are the constituents of cathode rays. (Or, at least, science still endorses the sentence 'electrons are the constituents of cathode rays' even though that sentence now means something quite different than it did for Thomson.) The basic theory accepted in Millikan's time has been rejected too, but we still hold the theory that 'electrons' can be attached to oil drops which then move in particular ways in electric fields as Millikan predicted they do. (Or again, science still endorses that sentence.)

Given this dependence on theory, albeit 'lower level' theory, entity realism is revealed as another attempt to respond to the challenge of theory-change, just like selective realism. It is in fact a version of what might be called *partial* realism. The basic, and intuitively appealing, idea behind any such view is to respond to the challenge of theory-change by abandoning realism about basic, fundamental theories and instead going down the theory-hierarchy until a lower level is found that has (allegedly) been unaffected by theory-change. So, the intuitive idea is that while, for example, our notions of atoms and of what happens during chemical combination are radically different from those that were around at the nineteenth-century, science still holds, as it did then, that, in appropriate circumstances, two volumes of hydrogen atoms plus one volume of oxygen atoms produces two volumes of water molecules. It follows that it is sensible (perhaps mandatory) to be realist about chemical molecules and their interactions, but not sensible to be realist about the basic constitutions of those molecules or about how those interactions occur – exactly because science's ideas about what that basic constitution is and how those interactions occur have undergone seemingly radical changes over time.

Notice that the partial realism involved is inevitably structuralist. The entity realist explicitly admits that she has no commitment to any view as to what electrons fundamentally are; she is committed to 'them' only *qua* those things, whatever they may fundamentally be, that are involved in the generalisations entailed by those lower-level theories to which she is committed.

Of course, a realist attitude toward successful scientific theories in general, including the highest-level ones, entails a realist attitude toward theories lower down the hierarchy, since the latter are logically weaker than the former. So, no one advocating a more robust version of scientific realism is going to attack realism about 'entities', once this has been identified as realism about lower-level theories that have arguably avoided change over time. The question is whether or not 'entity realism' gives up too easily on those higher-level theories that have *not* avoided theory-change. It is not clear whether Hacking recommends an outright anti-realist attitude toward top-level theories or is agnostic about them (he certainly seems to regard the

'endless disputes' about them as not worthy of attention). But since, in the end, it is the 'no miracles intuition' that underwrites a realist attitude even toward lower-level theories (Worrall, [2007], p. 146) and since it is top-level theories, like Fresnel's, that give us the most striking instances of predictive success that induce that intuition, it seems odd not to have something realist to say about them. Entity realism is somewhat confusingly expressed, but not, when judged from the viewpoint that I endorse, wrong; instead it is too weak: we should be able to underwrite a (no doubt somewhat sophisticated) realist attitude toward top-level theories as well as toward the lower-level theories that underwrite claims about 'manipulation of entities'. And SR claims that we can - as we have seen and as we shall soon see in more detail.

However, many commentators see SR itself as facing problems so severe that the position is effectively refuted. If that were correct then, given that Selective Realism seems indefensible, the conclusion beckons that there is no tenable version of scientific realism that endorses a realist attitude toward successful top-level or frontier theories; and so entity realism is, despite, indeed because of its relative weakness, the only defensible realist option. Fortunately, as I will argue in the next and final section of this paper, the problems that many have seen as discrediting SR can in fact readily be solved.

4 Criticisms of Structural Realism – Articulated and Countered

I will consider what I take to be the two most challenging problems that SR has been alleged to face. I will argue that the first is a genuine (though entirely foreseen) problem to which however there was always a ready-made solution. And the second problem, which is widely regarded as fatal to the position, is in fact easily met.²

4.1 The Fresnel- Maxwell example is entirely unrepresentative

One immediate criticism of SR (voiced, for example, by Colin Howson ([2000], p. 40)) is that the 'hackneyed' motivating historical example of the shift from Fresnel's classical wave theory of light to Maxwell's electromagnetic theory is entirely unrepresentative of theory-changes in the history of science. In fact, I conceded this already in my initial 1989 paper (pp. 120–21). The shift from Fresnel to Maxwell is unique amongst theory-changes history of science in being a case where the mathematical equations of the earlier theory (for example, Fresnel's equations for the intensities of the reflected and refracted beams in partial reflection, polarised in the plane of reflection and orthogonal to it) are reproduced *completely intact* within the later theory. (At any rate, the Fresnel-Maxwell case is the only one known to me, and the only one ever analysed in the philosophy of science literature, that involves complete retention of the older theory's mathematical equations within the newer theory). A view of the epistemic credentials of scientific theories that claims to respond to the *challenge* from theory-change in science, but in fact responds only to one single instance of theory-change, hardly seems to be on solid ground.

The position, then, that the mathematical structure of once successful theories is standardly entirely retained within their (even more successful) successor theories is entirely untenable – being directly at odds with every instance of theory-change in science, bar one. It does, however,

² This section borrows from and develops the points made in my [2007].

seem to a tenable view that in all such cases of theory-change (where, remember, the superseded theory has been genuinely predictively successful), the mathematical equations of the superseded theory (together of course with their connections to the phenomena) are 'substantially retained' within its successor theory, but *with corrections*. In many cases, the earlier and later theories satisfy what has been called in the literature 'the correspondence principle': the mathematical structure of the earlier theory re-emerges from that of the later theory, under some limiting process, where, moreover, the limit almost obtains (or better: actually obtains within experimental error) in the area in which the earlier theory enjoyed its predictive success.

The classic case is, of course, the relationship between the Special Theory of Relativity and its predecessor theory, Newtonian mechanics. The relativistic equations tend to the classical Newtonian equations as the ratio v^2/c^2 tends to zero, where v is the velocity of the body whose motion is under consideration and c is the velocity of light. This relationship means that the relativistic predictions about a body's motion and the corresponding classical ones, while strictly different, are entirely empirically indistinguishable for 'ordinary' motions, where the velocity involved is a tiny fraction of the velocity of light. (Standardly, the relativistic prediction multiplies the classical one by the factor $\sqrt{1 - v^2/c^2}$ which is practically indistinguishable from 1 in the case of 'ordinary' motions, where v is tiny compared to the velocity of light.) Notice that this means that, because it satisfies the correspondence principle, the Newton-Einstein case supplies, just like the Fresnel-Maxwell case, exactly the sort of explanation of the earlier theory's success that I described as crucial in defending the 'no miracles intuition' against theory-change and hence in underwriting a tenable form of scientific realism. Just as Maxwell's theory explains the empirical success of Fresnel's (as 'non-miraculous') by yielding exactly the same mathematical equations (which are of course tied to observational results in the same ways); so, differently but equally, the Special Theory of Relativity explains the empirical success of Newton's by yielding mathematical equations which, while strictly different, are nonetheless empirically indistinguishable from those yielded by Newton's theory for the domain (of relatively slowly moving objects, such as planets) in which Newton's theory had enjoyed its predictive success. No wonder then, from the vantage point of Special Relativity, that Newtonian predictions – about for example the existence of the hitherto-unsuspected planet, Neptune – were correct.

So, SR is based on the claim that the history of predictively successful science has been one of structural continuity *modulo* the correspondence principle. Given that claim, SR is then, if you like, one fairly innocuous (optimistic) inductive step away: the step being to the conclusion that our current theories are likely to be approximately true in the (very weak) sense that, whatever theories may eventually replace them, our current theories will be retained within those new theories in the same way that previous theories were retained within them, almost certainly courtesy of some – of course currently unspecifiable – application of the correspondence principle. This may be too weak a notion of continuity for some, but it seems to be the strongest notion that is compatible with the history of theory-change in science. Correspondingly, the structural realist position based on that alleged continuity may, correspondingly, seem to many too weak a view to count as 'real' realism. If so, that's fine: call it 'Structural Quasi-Realism' if you like, or even 'Structural Empiricism' – names of course do not matter. It does in fact seem appropriate to regard it as a version of realism because it respects the 'no miracles intuition': there is something correct about our current theories (beyond their directly empirical consequences) that will be retained, albeit in modified form, within later theories and which accounts for the predictive success of those theories. The feeling that a view needs something

stronger if it is to count as ‘real’ realism seems to be based on (a) a wish that the history of theory change was simpler, more straightforwardly cumulative than it in fact is and (b) the feeling that what a realist really needs to claim is that we have some non-theory-mediated access to the furniture of the universe and some reason to claim that our theories correspond with that reality. But everyone agrees that this feeling, once articulated, is entirely fanciful (though it nonetheless often seems to be presupposed as an unarticulated assumption). If that is what it takes to defend realism, then realism is of course indefensible. So long as we restrict ourselves to tenable, non-fanciful positions, SR is as realist as it gets.

4.2 ‘The Newman Objection’ (or rather ‘The Newman “Objection”’)

The problem that is most widely regarded as fatal to the prospects of SR is posed by the so-called ‘Newman objection’. A little background is required to understand the issues involved.

SR accepts that the only tenable account of the meaning of theoretical terms is the descriptive account. If asked what the term ‘electron’ means according to current lights, then we can only recite the (full) current theory of electrons: electrons, if there are any, are the entities – whatever they are – that do the things attributed to them in that theory. And the term ‘electron’ (as currently construed) refers if and only if our current theory is true – literally true, not approximately true, whatever that means. Since SR accepts that it would be foolish to hold that our current theories are literally true, it is, as noted earlier, ‘realism without (accredited) reference’.

The only (alleged) alternative account to the descriptive one is in effect Putnam’s causal theory of reference already mentioned in section three. Putnam’s treatment may seem plausible as an account of the meaning of terms for ordinary everyday objects (of which we think of ourselves as having knowledge ‘by acquaintance’ (in Russell’s terminology)). But it is, I claim, entirely untenable as an account of the meaning/reference of *theoretical* terms. The best way to see this is through Hacking’s analysis in *Representing and Intervening* ([1983], ch. 6), which constitutes, ironically, a valiant attempt to defend Putnam’s account. As Hacking makes clear, Putnam’s account of the meaning of a term involves (amongst other things like a ‘stereotype’ – irrelevant to genuinely theoretical terms like ‘electron’) the *extension* of the term. Hacking indicates this by ‘...’ – the ‘dots of extension’. But these dots are mystical: we cannot point to the extension of the term electrons, nor, as Hacking admits in dealing with his own example (of ‘glyptodon’), can we ‘put them on the page of a dictionary’ (*ibid.*, p. 80) The only way we can get at the extension (which *is* the reference!) is via our theories. The only way to avoid the conclusion that Putnam’s ‘causal’ account collapses into the descriptive theory of meaning is via our old friend the assumption that we can have some non-theory-mediated access to the basic constituents of the world, in this case electrons, so that we can compare what we say about them with the ‘real things’. But as noted, although strongly encouraged by the way that semantics is taught in logic courses, this assumption once articulated is clearly untenable. We are left (and should be happy to be left) with the descriptive theory of the meaning/reference of theoretical terms.

The descriptive account has the clear and immediate implication that the cognitive content of any theory T is fully captured by its Ramsey Sentence, $R(T)$. $R(T)$, as is well-known, is constructed by replacing all the *theoretical* predicates in T by second-order variables and then

existentially quantifying over those variables.³ The claim that $R(T)$ captures the full cognitive content of T simply, then, reflects the fact that, so far as our fundamental, primitive theoretical notions are concerned, we know about them only by description – that is, via their role in our theories. If asked what an electron, say, is (or rather what we *think* an electron is), one can do no more than recite our current (full!) theory of electrons – electrons (if they exist at all) are *whatever it is* that satisfy our current relevant theories. But this means that a theoretical term like ‘electron’ in effect plays the role of a (second-order) ‘ambiguous name’; and, as is well-known, in systems of logic that employ ambiguous names α_i , $P(\alpha_i)$ is logically equivalent to $\exists x P(x)$ (See e.g. Suppes [1957]). (Notice that the account I defend takes the Ramsey sentence of T in its original Ramsey form, rather than its Lewisian variant⁴ (Lewis [1970])). This is, after all, exactly what Structural Realism asserts: we have (tentative) knowledge of, tentative access to, theoretical ‘entities’ only through the way that they (together of course with the other primitive theoretical predicates involved in our best theories) structure the phenomena.

However, as we just noted, it is widely believed that any account that claims that the cognitive content of any theory is the same as that of its Ramsey sentence has been shown to be entirely and irredeemably untenable. The argument at issue was first presented in a 1928 paper by the Cambridge mathematician Max Newman, responding to Russell’s version of SR. Newman’s argument was brought back to the attention of philosophers of science via a 1985 article by Demopoulos and Friedman. The argument in its crispest form goes as follows:

1. SR is, as I just re-acknowledged, committed to the view that the Ramsey sentence of any scientific theory T captures the full ‘cognitive content’ of that theory.
2. However, as Newman showed, the Ramsey sentence of any theory imposes only a very weak constraint on the universe – it amounts in essence to a mere cardinality constraint, and so if there are sufficiently many objects in the universe, then $R(T)$, for any T , will automatically be true.
3. However, as is obvious, the truth of standard scientific theories requires much more from the universe than that it have an appropriate number of elements.
4. Hence SR is committed to an account of the cognitive content of scientific theories that is plainly untenable and is, therefore, itself untenable.

No wonder, then, Russell immediately gave up on SR in the light of Newman’s result:

Russell quickly abandoned SR when Newman showed that any set with the right cardinality could be arranged so as to have the same structure as the world – a result analogous to that claimed in Putnam’s model-theoretic argument against realist theories of reference. (Griffin [1998], p. 400)

³ Of course, if you presuppose set theory then the theoretical predicates can be thought of as varying over sets and then the Ramsey sentence is entirely first-order.

⁴ Lewis’s argument for reading the quantifiers as ‘there are a *unique* \emptyset_i ’s such that ...’ seems to me entirely unconvincing. To see why, indulge in a little thought experiment. Suppose tomorrow we “discover” (i.e. produce a new and massively confirmed theory that entails) that there are two distinct types of elementary particles with unit electrical charge. One type let’s call them electrons₁ do all (and only) the things that our current theory of electrons asserts that they do; electrons₂ also do all those things but have further properties and effects not shared by electrons₁. From the vantage point of the later theory, the Lewis sentence for our current theory of electrons is false. But that is surely the wrong judgment: the earlier theory was really true but incomplete – and this judgement is delivered by considering the content of the theory as capture by its real, not-Lewisian Ramsey sentence.

Of course, this leaves open the possibility that there is something especially faulty with Russell's version of SR, but in all respects relevant to the current discussion this is not true (or at any rate I do not believe it to be true). If so, then advocates of SR such as myself and Elie Zahar ([2001]) seem to have shown reprehensible ignorance of the literature in advocating a position that had already been conclusively demolished before we published.

In fact, the 'Newman Objection' is no sort of objection at all to SR. Before seeing exactly why, reflect on how surprising it would be if Newman's argument, as is very widely believed, were cogent. The argument claims to refute not just SR but any position that accepts that the full epistemic content of any theory is captured by its Ramsey sentence. This is because the argument claims to show that that Ramsey sentence places only a cardinality constraint on the universe (that is, that all that the Ramsey sentence of any theory requires of the universe in order for that sentence to be true is that the universe includes sufficiently many individuals.) But the view that its Ramsey sentence captures the full cognitive content of any theory T , is, as we saw, a direct consequence of accepting the descriptive account of the meaning of theoretical terms. Since that account is surely the only defensible one, then the Newman argument, if cogent, would entail that all human theorising does no more than impose a constraint on how many individuals there are in the universe – a *reductio ad absurdum*, if ever there was one!

But Demopoulos and Friedman, in reintroducing the Newman Objection to the realism debate, seem to assert exactly that Newman established that $R(T)$ for any T can do no more than impose such a cardinality constraint. They write:

The difficulty is with the claim that *only* structure [as revealed by the Ramsey sentence] is known. On this view 'the world consists of objects, forming an aggregate whose structure with regard to a certain relation R is known, say [it has] structure W ; but of R nothing is known ... but its existence; ... [A]ll we can say is *There is a relation R such that the structure of the external world with reference to R is W .*' But '*any* collection of things can be organized so as to have the structure W , provided there are the right number of them'. Thus, on this view, only cardinality questions are open to discovery! Every other claim about the world that can be known at all can be known a priori as a logical consequence of the existence of a set of α -many objects. For, any given set A of cardinality α , can with a minimum of set theory or second-order logic establish the existence of a relation having structure W , provided that W is compatible with the cardinality constraint that $|A| = \alpha$. (The relevant theorem from set theory or second-order logic is the proposition that every set A determines a full structure, that is, one which contains every relation [in extension] of every arity on A ...). (*op. cit.*, pp. 627–28)

The crucial mistake made by Newman (and also, at any rate in the passage just quoted, by Demopoulos and Friedman) is to assume that in order to identify the structural claims of a theory T , that is, in constructing T 's Ramsey sentence $R(T)$, one should replace *all* the predicates occurring in T by predicate variables and existentially quantify over them – in other words, *all* the predicates in the theory should be treated as theoretical.⁵ It is surprising that Russell should have conceded to Newman (as both the *Routledge Encyclopedia* and Demopoulos and Friedman correctly report that he did), given that his whole structural realist view was based (and

⁵ This section of the paper follows the treatment in J. Worrall and E. Zahar, 'Appendix IV: Ramseyfication and Structural Realism' in Zahar ([2001]), pp. 236-251. I am in general greatly indebted to Elie Zahar for many long discussions of structural realism.

necessarily based) on a sharp distinction between theoretical and observational predicates (in his own terms: between things of which we can have knowledge *by acquaintance* and others of which we can have knowledge only *by description*).

It is true that if *all* predicates are quantified-over, then the resulting Ramsey sentence is hopelessly weak and so could not possibly be construed as capturing the full cognitive content of the theory from which it is constructed. (Even then, however, the situation, as we'll shortly see, is a good deal more complicated than Demopoulos and Friedman, following Newman, assert.) But the real Ramsey sentence *turns only theoretical and not observational predicates* into variables and then existentially quantifies over them. Given that some predicates are to be taken as observational and therefore as interpreted independently of theory and so emphatically *not* to be quantified-over, the charge that the Ramsey sentence of any theory imposes unacceptably weak constraints on the universe immediately dissolves. It is in fact well-known that, with the observational predicates fixed, the original theory T and its Ramsey sentence $R(T)$ are empirically equally powerful: that is, every sentence that is expressible purely in observational terms and is deducible from T is also deducible from $R(T)$. Hence $R(T)$ is co-refutable with, has the same empirical content as, T ; and so, assuming that T itself has non-empty empirical content, its Ramsey sentence, $R(T)$, far from imposing merely a cardinality constraint on the universe (or otherwise being obviously too weak), in fact requires all the same phenomena to hold as does the original T .

If we suppose – as Newman himself did, albeit incorrectly, as we just saw – that there is no distinction to be made between observational and theoretical predicates in T , and so all predicates are to be replaced by variables and quantified over, then that (aberrant) ‘Ramsey sentence’ formed in this way from T (let’s call it $R^*(T)$) will indeed be a very weak claim – being in effect a purely mathematical statement. However even then Newman’s claim (echoed by Demopoulos and Friedman) requires modification. Newman suggested that if structuralism were correct, that is, the Ramsey sentence was all that could be known, then this would entail that all that can be known ‘theoretically’ is ‘the number of constituting objects’ (Newman [1928], p. 144). But in fact, only in particular cases does the sentence, $R^*(T)$, formed from a theory T by quantifying over *all* of T ’s predicates (whether theoretical *or* observational), determine the cardinality of the domains that satisfy that theory. For example, if T is a contradiction then so is $R^*(T)$; while if T is $\forall x(S(x) \leftrightarrow \neg(x=x))$ then $R^*(T)$ is $\exists \Phi (\forall x(\Phi(x) \leftrightarrow \neg(x=x)))$ which is a logical truth, since $\Phi(x)$ can be chosen to be the predicate $\neg(x=x)$. In the first case, there are no models and in the second case, all interpretations are models. In neither case is there, of course, any restriction on cardinality. Moreover, by the upwards and downwards Löwenheim-Skølem theorems, if T is any theory with an infinite model of any cardinality then it has models of *all* infinite cardinalities. Hence the Ramsey sentence of any logically true theory T will be true in virtue of a range of structures with domains of all infinite cardinalities. It is only in the case of a theory that has only finite models that its Ramsey sentence determines the size of the domain of individuals. So, for example, if our theory were $\forall x Sx \ \& \ \exists x_1, x_2, \dots, x_n (\neg(x_1 = x_2)) \ \& \dots \ \& \ \neg(x_{n-1} = x_n) \ \& \ \forall x(Sx \leftrightarrow (x = x_1) \vee \dots \vee (x = x_n))$ then its Ramsey sentence would be equivalent to the assertion that the domain of individuals consists of exactly n members.

Demopoulos and Friedman, as we saw, went along with Newman’s mistaken assumption that all predicates are to be made into variables and quantified over in the passage I just cited (*op. cit.*, pp. 627–28). However, they do elsewhere in the same article state:

[...] if our theory is consistent and *if all its purely observational consequences are true*, then the truth of the Ramsey-sentence follows as a theorem of set-theory or second-order logic, provided our initial domain has the right cardinality... (*op. cit.*, p. 635)

And perhaps this is their real position. It is, of course, significantly different from Newman's original position and correspondingly not open to the above objection. The claim now seems to be that although the Ramsey sentence of any theory T says everything that T does about observational matters, it is nevertheless somehow not strong enough to capture the full cognitive content of T .

It is difficult to know what precisely to make of this claim. Notice that the empirical equivalence of T and its Ramsey sentence $R(T)$ extends far beyond what we might normally take that phrase to mean – it extends far beyond the fact that T and $R(T)$ 'save the same phenomena'. The Copernican and Ptolemaic theories, for example, are (or can be made) empirically equivalent with respect to observed planetary motions (courtesy of delicate adjustment of epicycles); but any theory T and its Ramsey sentence, $R(T)$, are observationally equivalent in the much stronger sense that *any sentence* that is expressed in the observational vocabulary and is entailed by T is also entailed by its Ramsey sentence. Many sentences expressed in purely observational vocabulary should count, in any one's book, as theoretical: one often cited example is the claim that there are unobservables – that is, individuals possessing no observable property.

Moreover, the Ramsey sentence of any theory is itself observational in the sense that it is expressed purely in observational language – the only named predicates in it being observational. This makes the claim just quoted from Demopoulos and Friedman – that '[...] if our theory is consistent and *if all its purely observational consequences are true*, then the truth of the Ramsey-sentence follows as a theorem of set-theory or second-order logic, provided our initial domain has the right cardinality' – in a sense trivially true, even without the cardinality proviso. This is because (a) T 's Ramsey sentence $R(T)$ is provably a logical consequence (of course second-order logical consequence) of T and (b) $R(T)$, like any other sentence, entails itself.

Those who have assumed that even this qualified result of Demopoulos and Friedman spells the end for SR, at least as version of realism, have perhaps read their claim as asserting that if a theory T 'is true at the phenomenal level' then its Ramsey sentence $R(T)$ is automatically true, full stop. But if 'T is true at the phenomenal level' means that all its *singular*, observationally decidable sentences are true, then the truth of the Ramsey sentence $R(T)$ does *not* follow, as can readily be shown using the compactness theorem (*op. cit.*, p. 245).

Even before Newman's and Demopoulos and Friedman's article, it was often assumed that the Ramsey sentence of T eliminates the theoretical content of T along with its theoretical predicates. But quantifying over the erstwhile theoretical predicates whilst leaving the observational predicates fixed, removes those theoretical predicates merely linguistically, and the theoretical notions surely live on within the Ramsey sentence via the structure that they impose on the observational content. That is, after all, exactly what structural realism is about – far from eliminating theoretical entities, it merely asserts that our sole access to them is via the structure that they (jointly) impose on observation. And if it asserts that we have access to them then it asserts that they exist (dependent of course on the truth of the theory). Indeed, as Elie Zahar pointed out to me, if we follow Quine's dictum that 'to be is to be quantified over' (or, better: 'to be asserted is to be quantified over in a sentence that you assert') then the Ramsey sentence of some theory $T(S_1, \dots, S_n, O_1, \dots, O_r)$, where the S_i are theoretical predicates and the O_j

are observational, namely $\exists\Phi_1 \dots \exists\Phi_n (T(\Phi_1, \dots, \Phi_n, O_1, \dots, O_r))$ clearly asserts that the ‘natural kinds’ S_1, \dots, S_n (the extensions of the theoretical predicates S_1, \dots, S_n in the initial theory T) exist in reality just as realists want to say. It is just that – as always – we fool ourselves if we think that we have any independent grip on what the S_1, \dots, S_n are aside from whatever satisfies the relevant Ramsey sentence (assuming that the theory whose Ramsey sentence we are considering is true).

In the end, then, the much-vaunted ‘Newman objection’ is no real objection to SR at all. In its original form in Newman’s own paper, it relies on a false assumption (the assumption that SR treats all predicates in any scientific theory as theoretical); and, in its modified form in the article by Demopoulos and Friedman it simply highlights a consequence of SR that is essential to it. The fact that the ‘objection’ has been found so troubling depends, I think, on nothing more than a residual vague feeling that if this is what SR amounts to then it cannot count as ‘*real* realism’ – a charge that I already rejected earlier.

Finally notice that the claim, essential to SR, that the full cognitive content of any theory is captured by its Ramsey sentence further underwrites the claim that SR is not a version of selective realism. The mathematical structure of Fresnel’s classical wave theory *is given by* its Ramsey sentence. So, the fact that the mathematical structure of Fresnel’s theory is recovered within Maxwell’s theory means that the *whole of the cognitive content* of that theory is retained – not a part of it. All that is rejected are a few elusive bells and whistles of a metaphysical kind that encourage various mental ‘models’ of the light-carrying medium, such as an interconnected set of weights on springs, which can be pedagogically useful but are entirely unnecessary and often misleading and not significant parts of the scientific content of the theory. As Poincaré put it, remember, (*op. cit.*, p. 8), all that is changed in the switch from Fresnel to Maxwell are ‘merely names of the images which we substituted for the real objects which Nature will hide for ever from our eyes.’

5 Conclusion

The only defensible realist reaction to the challenge posed by the history of theory-change is, I have argued, the ‘type two reaction’: earlier (predictively successful theories), although now rejected, continue to somehow be reflected in later (even more successful theories). The attempts by SR’s main rivals – Selective Realism and Entity Realism – to provide a defensible form of this reaction are, for different reasons, unacceptable. SR, on the other hand, despite various criticisms that some have incorrectly regarded as fatal, stands as a viable reaction to the challenge of theory-change. While it might not seem ‘realist enough’ for some tastes, it is as strong a version of realism as is compatible with the facts about theory-change. It appears to be the only viable scientific realist game in town.

References

Carman, C. and Díez, J. [2015]: ‘Did Ptolemy Make Novel Predictions? Launching Ptolemaic Astronomy into the Scientific Realism Debate’, *Studies in History and Philosophy of Science*, **52**, pp. 20–34.

- Demopoulos, W. and Friedman, M. [1985]: 'Critical Notice: Bertrand Russell's *The Analysis of Matter*: Its Historical Context and Contemporary Interest', *Philosophy of Science*, **52** (4), pp. 621-39.
- Fahrbach, L. [2011]: 'How the Growth of Science Ends Theory Change', *Synthese*, **180**, pp. 139-55.
- Fahrbach, L. [2017]: 'Scientific Revolutions and the Explosion of Scientific Evidence', *Synthese*, **194**, pp. 5039-72.
- Fresnel, A. [1818]: 'Memoire sur la Diffraction de la Lumière (Coronné par l'Academie des Sciences)', in Fresnel, A., *Oeuvres*, vol. I, pp. 247-382.
- Griffin, N. [1998]: 'Russell, Bertrand Arthur William (1872-1970)', in Craig, E. (ed.), *Routledge Encyclopedia of Philosophy*, London: Routledge.
- Hacking, I. [1983]: *Representing and Intervening: Introductory Topics in the Philosophy of Natural Science*, Cambridge: Cambridge University Press.
- Hacking, I. [1988]: 'The Participant Irrealist At Large in the Laboratory', *The British Journal for the Philosophy of Science*, **39**, pp. 277-94.
- Hardin, C. L. and Rosenberg, A. [1982]: 'In Defense of Convergent Realism', *Philosophy of Science*, **49**, pp. 604-15.
- Howson, C. [2000]: *Hume's Problem: Induction and the Justification of Belief*, Oxford: Oxford University Press.
- Kitcher, P. [1993]: *The Advancement of Science: Science without Legend, Objectivity without Illusion*, Oxford: Oxford University Press.
- Kitcher, P. [2001]: 'Real Realism: The Galilean Strategy', *The Philosophical Review*, **110**, pp. 151-97.
- Ladyman, J. [2011]: 'Structural Realism Versus Standard Scientific Realism: The Case of Phlogiston and Dephlogisticated Air', *Synthese*, **180**, pp. 87-101.
- Laudan, L. [1981]: 'A Confutation of Convergent Realism', *Philosophy of Science*, **48**, pp. 19-49.
- Lewis, D. [1970]: 'How to Define Theoretical Terms', *Journal of Philosophy*, **67** (13), pp. 427-46.

- Lipton, P. [2000]: 'Tracking Track Records', *Proceedings of the Aristotelian Society*, Supplementary Volume LXXIV.
- Newman, M. H. A. [1928]: 'Mr. Russell's "Causal Theory of Perception"', *Mind*, **37** (146), pp. 137–48.
- Newton, I. [1934]: *Sir Isaac Newton's Mathematical Principles of Natural Philosophy and His System of the World*. Trans. Motte, A., Revised by Cajori, F, Berkeley: University of California Press.
- Papineau, D. and Cruse, P. [2002]: 'Scientific Realism Without Reference', in Marsonet, M. (ed.), *The Problem of Realism*, Aldershot: Ashgate.
- Poincaré, H. [1952]: *Science and Hypothesis*, New York: Dover.
- Psillos, S. [1991]: *Scientific Realism: How Science Tracks Truth*, London: Routledge.
- Putnam, H. [1975a]: 'Mathematics, Matter and Method', in *Philosophical Papers*, Vol. 1, Cambridge: Cambridge University Press.
- Putnam, H. [1975b]: 'Mind, Language, and Reality', in *Philosophical Papers*, Vol. 2, Cambridge: Cambridge University Press.
- Suppes, P. [1957]: *Introduction to Logic*, New York: Van Nostrand.
- Worrall, J. [1989]: 'Structural Realism: The Best of Both Worlds?', *Dialectica*, **43**, pp. 99–124.
- Worrall, J. [2007]: 'Miracles and Models: Why Reports of The Death of Structural Realism May Be Exaggerated', *Royal Institute of Philosophy Supplements*, **82** (61), pp. 125-54
- Zahar, E. [2001]: *Poincare's Philosophy: From Conventionalism to Phenomenology*, Chicago: Open Court.