

Replacing Recipe Realism

Juha Saatsi
University of Leeds
J.T.Saatsi@leeds.ac.uk

26th October 2015

Abstract

Many realist writings exemplify the spirit of ‘recipe realism’. Here I characterise recipe realism, challenge it, and propose replacing it with ‘exemplar realism’. This alternative understanding of realism is more piecemeal, robust, and better in tune with scientists’ own attitude towards their best theories, and thus to be preferred.

1 Introduction

Alice: *Hofstadter’s 40-year old prediction of the butterfly spectrum of electrons energy in a magnetic field was just confirmed!*

Bob: *Wow! Realism about quantum mechanics further vindicated!*

Alice: *But we’ve been here before, no? The possibility of impressive predictions from radically mistaken premises. Remember the ether models of light. Underdetermination, and all that!*

Bob: *Well, I am a structural realist! There is a structural correspondence between successful ether models and Maxwell’s theory, and probably one in any actual example of underdetermination.*

Alice: *So, given any predictively successful, mature theory, you bet it gets ‘the structure’ right?*

Bob: *Yup. And by ‘structure’ I mean... Well, take the Ramsey-sentence of a theory, and ...*

Conversations of this kind have echoed widely in the recent philosophy of science. Never mind Bob’s structural realist stripe. What should strike you worrisome is how readily Bob attests to particular realist commitments, without any need to consider the further details concerning the theoretical prediction in question. It is as if Bob had an abstract realist recipe in his

sleeve, capable of distilling the trustworthy aspects of a theory, applicable to *any* good, predictively successful mature theory.

Bob is a *recipe realist*. Many contemporary realists are like Bob: they readily declare their epistemological commitments in the abstract, recipe like, regarding all/most/many theories in mature science, on the basis of few detailed exemplars that are meant to motivate and illustrate their favourite recipe. But even the most closely studied exemplars are open to a wide range of interpretation; consequently there are almost as many recipes as there are recipe realists. Understandably there are quarrels over the best recipe, but recipe realists—unified in their spirit—vehemently oppose anti-realists who claim that recipe realism should be given up altogether, in favour of empiricism or instrumentalism.

I have no truck with the anti-realist side of the debate, but I think recipe realism is deeply flawed. It is flawed in its spirit, not just in the letter of this or that particular recipe. In this paper I aim to characterise recipe realism, challenge it, and propose replacing it with (what I here call) *exemplar realism*. This alternative understanding of realism is more piecemeal, robust, and better in tune with scientists' own attitude towards their best theories, and thus to be preferred.¹

2 Debates about recipe realism

There is a battery of familiar challenges to scientific realism regarding well-confirmed theories that are thoroughly accepted by scientists. Some turn on evidence from the history of science against the realist notion that we can have good scientific evidence for the 'approximate' or 'partial' truth of our current best theories (or something like that). This challenge effectively arises out of inconsistency between past and current science. Other (perhaps less familiar) challenges arise out of other inconsistencies in science: internally inconsistent models, for example, or the indispensable use of idealisations that are radically at odds with our background beliefs.

Realist responses to these challenges are equally familiar. According to these responses the inconsistencies above are compatible with the kind of continuity or correspondence that a suitably modest realism requires. The intended sense of continuity or correspondence is signalled by the familiar realist monikers (such as 'structural realism' or 'entity realism'). Specifying the precise sense of continuity or correspondence is tricky, and there is much divergence in the details of different realist responses. For example, the sense of structural continuity associated with dispositional semi-realism is quite different from that of Ramsey-sentence structuralism. But while the various

¹This 'big picture' paper can merely sketch this alternative vision of realism, leaving many of the details to be developed elsewhere. See Saatsi (forthcoming a, b) for related discussion.

realist responses are multifarious in their letter, they are broadly unified in their spirit: by and large, realists have aimed to provide a recipe for characterising the truth-content of our current theories (or, more broadly, characterising a uniform sense in which well-confirmed theories ‘latch onto’ unobservable reality).

There is much debate within the realist camp regarding, say, the most appropriate formal framework for characterising realist commitments. There are broad-ranging debates about how to best understand the notion of ‘structure’, for example, and how it differs from the ‘nature’ of unobservable entities (if at all). But these debates regarding the letter of the realist response should not hide the considerable underlying unity in the spirit of the overall realist endeavour: realists by and large agree that it is meaningful and worthwhile to aim to delineate our epistemic commitments in the abstract, largely independently of case-specific details, resulting in a recipe (or two) projectable both to future science and well beyond the selected case-studies that arguably fit the recipe.

The spirit of recipe realism is prevalent in the realism debate, on both sides of the debate. It shows in the way the realists and anti-realists alike go looking for cases from the history of science to test (or ‘falsify’) particular realist recipes, for example. For instance, Stanford (2003) argues in this way against Worrall’s structural realism on the basis of a problematic case from the history of biology: arguably the mathematical structure expressed by Galton’s Ancestral Law of his “stirp” theory of inheritance—the driver of the theory’s predictive success—is not preserved in contemporary genetics. Theories of inheritance bear little resemblance to the cases (from physics) that have motivated structural realism in the first place, but Stanford’s dialectic makes perfect sense if we construe structural realism as purportedly offering a general recipe that applies across the sciences. And—to be fair—Worrall’s (1989) discussion of structural realism is fully general in its tone, despite primarily drawing motivation from ‘revolutionary’ theory shifts in physics. In Worrall’s words, structural realism aims to ‘to give the [anti-realist] argument from scientific revolutions its full weight and yet still adopt some sort of realist attitude towards presently accepted theories in physics and *elsewhere*.’ (p. 99, my emphasis) Similarly, later structural realists have presented their epistemic credo (‘all the we know is structure’) in very general terms (see e.g. French, 2014; Ladyman et al. 2007). Structural realists have advocated recipe realism, and amongst these realists the debate has mostly revolved around the letter of the best structuralist recipe. Is the recipe that delineates the structuralist’s epistemic commitments in the abstract best expressed in terms of a theory’s Ramsey-sentence? In terms of partial structures? In modal terms? Anti-realists like Stanford, on the other hand, who have grown skeptical of the whole recipe realist endeavour, see the ‘falsification’ of this or that realist recipe as a motivation for abandoning realism altogether in favour of neo-instrumentalism, say. For if it looks like

there are ‘counterexamples’ to be found to any realist-recipe, what becomes of the hope to find a recipe worthy of realist commitment?

Some of the most recent realist writings also exemplify the spirit of recipe realism.² Consider first Vickers (2013) who is concerned with particular derivations of successful predictions. What epistemic attitude should we adopt towards derivations of successful predictions, given that the history of science shows (in various ways) how successful predictions can be derived from clearly mistaken premises? In response to this question Vickers proposes ‘a new “recipe” for the prospective identification of (at least some) working/idle posits’ (p. 189). After some twists and turns Vickers ultimately arrives at a modest conclusion that nevertheless expresses the spirit of recipe realism: ‘Even if the “working posits” of contemporary science cannot be prospectively identified, it remains possible that we might develop a recipe for identifying certain idle posits.’ (p. 209)

Or consider Peters (2014) who is concerned not with particular derivations of successful predictions, but with the predictive success of a given (unifying) theory. Peters first sensibly distances his realist brief from the kinds of recipes (like structural realism) that ‘regard a particular type of theoretical element [such as ‘mathematical structure’] as having a special explanatory role’ for the realist.³ But then, in offering his own ‘unification account’, Peters still ambitiously aims to provide a general, prospectively applicable recipe—Peters calls it ‘algorithm’—for identifying the elements of successful scientific theories that we can selectively be realists about. Roughly speaking, the recipe that Peters has in mind with his ‘empirically successful subtheory account’ turns on the idea that empirical success appropriately understood requires a certain kind of unification that ‘does not justify regarding entire theories as confirmed but rather provides a [recipe] for picking out the confirmed elements within theories’ (p. 390).

Stepping back from these examples of recipe realism, we can characterise the overall realist endeavour exemplified therein as aiming to pin down a perspective on science that is sweeping and monolithic in its realist outlook. The aim is to capture theories’ epistemic commitments across a wide range of disciplines and different areas of scientific theorising in unified terms, offering

²Since the details of the positions reported here do not matter for my argument I will leave them aside.

³To this extent Peters’s comments are a very much in line with my criticism of recipe realism below (§3).

The first major objection to special accounts is that they are motivated by particular cases or types of cases that are not necessarily representative. We might concede that the mathematical structural realist and the phenomenological realist have successfully identified those theoretical elements that are intuitively essential to the Fresnel wave theory of light and the London model of superconductivity, respectively. And yet we can, and should, deny that similar analyses can be applied more generally. (p. 382)

recipes or algorithms that are first motivated by particular considerations and case-studies, and then optimistically projected well beyond those to the rest of science. Those who prefer a recipe of a different flavour can challenge those projections on the basis of this or that case that does not seem to fit the recipe on offer. Anti-realists go further in taking the seeming proliferation of increasingly weaker realist recipes to signal the need to give up on realism altogether in favour of instrumentalism (e.g. Stanford 2003).

Before I go on to challenge the spirit of recipe realism it is worth pausing to note that the prevalence of this spirit in the realism debate is far from surprising in the light of the debate's intellectual history. There is a long tradition to framing scientific realism in very global terms, as a philosophical theory about 'mature science'. As Putnam (1979) had it, for example, paraphrasing Boyd (1973), realism is a broad empirical hypothesis according to which 'terms in a mature science typically refer', and 'the laws of a theory belonging to a mature science are typically approximately true'. (p. 290) Over the 1980s and 90s, largely in reaction to Laudan's historical attack of realism, there has been gradual weakening of the sense of 'approximation' at stake. But realists have by and large maintained their global aspirations and the idea that realism is an explanatory theory about 'mature science'.

3 Challenging recipe realism

I hold a broadly realist outlook on science, but I also agree with those who deem the debates on recipe realist stale and degenerative. Instead of engaging those debates we can challenge the spirit of recipe realism altogether, making room for a realist endeavour of a different sort. Here are some reasons for doing this.

Diversity of science. In the face of the impressive diversity of science it is hard to think of a *prima facie* reason to expect that one or two realist recipes could fit all (or even most) of predictively successful mature science. Scientific theorising on the whole is rather inhomogeneous in its key aspects vis-à-vis the realism issue. Have a look at the kinds of explanations that are accepted in different areas of science, from particle cosmology to molecular biology, from systems biology to geology, and so on. Have a look at the kinds of theoretical reasoning that are successfully employed in connection with different research traditions, from gauge theories in fundamental physics to immunology, from climate science to evolutionary biology. Even within a single discipline one can find radically different modes of explanation and theoretical reasoning in play. Why think that we are apt to latch onto reality in the same way throughout the sciences, or even within a single discipline? Quite plausibly, some areas of science are more likely

to exhibit underdetermination than others.⁴ Some subject matters are very far removed from everyday reality (e.g. quantum fields), while others are relative close to it despite being about thoroughly unobservable entities and processes (e.g. causal-mechanistic systems in microbiology). In the face of all the diversity, why think that one (or even a handful) of recipes uniformly and fairly captures—across the board—the way in which theories’ empirical success is correlated with the way they latch onto reality?

Recipe realists are prone to abstracting away from real-life differences in science—in how explanatory and other kinds of theoretical virtues are judged, for example—in motivating their recipes. At a high enough level of abstraction all scientific theorising looks homogeneous, as when abstractly described as simply relying on the method of ‘inference to the best explanation’, the reliability of which has then become a central bone of contention. The realist idea that we can thus argue for wholesale realism, as an abductively justifiable theory about all of mature science, was a bad one. In the light of the various criticisms of global justification of inference to the best explanation through the No-Miracles Argument, there is no reason to think that the intuition underlying the argument, motivating a broad realist outlook on science, would thus undergird recipe realism. (See in particular Saatsi (forthcoming a) and Frost-Arnold, 2010.)

Varieties of realist explanations. One need not advocate recipe realism in order respect the intuition that the overall empirical success of science is (by and large) due to—and hence in some sense explained by—our theories latching better and better onto reality. This is because this intuition is respected even better with *pluralism* about the different ways in which a theory can be at the same time radically mistaken in some of its assumptions, yet empirically successful by virtue of suitably ‘latching onto reality’. The various illustrations of different realist recipes function as exemplars of this, of course, showing how empirical success can be due to getting ‘structure’ right, getting right some of a system’s less-specific features, getting right the fundamental causal contributors, and so on. Instead of thinking that different recipe proposals are in competition for the ‘one true recipe’, think of them as exploring this possibility space. Each recipe may fit a limited number of specific cases, furnishing realist explanations of empirical success (with respect to specific instances of predictive success, say) to the tune of the ‘no-miracles’ intuition. But all that of course provides no reason to think that those recipes should be projected across the board, and that the other possibilities are not realized somewhere else in science.

⁴See e.g. Smeenk (2013) and Butterfield (2012) on underdetermination in cosmology, and Belot (2014) on underdetermination in geology.

Contentless recipes. In effort to accommodate a variety of cases a recipe realist can describe her recipe in increasingly abstract and general terms. But the more abstract and generic a recipe becomes, the less work it can do for the realist in providing a clear sense of how we are meant to regard our theories latching onto reality. Stanford (2003) rightly complains that all realist attempts to pin down a general recipe seem to have resulted in a characterisation of epistemic commitments so vague and ill-defined as to deflate the realists claim to be able to trust in our current theories in the sense given by the recipe. For Stanford this challenge is (part of) his ‘trust argument’: in presenting a recipe for stating their epistemic commitment in a way that is compatible with various different kinds of theory-shifts in the history of science the realist is forced to ‘defend realist inferential entitlements that are so weak as to capitulate to the realist’s opponent on the question of whether we can safely trust the accounts of nature given by current or future successful scientific theories’ (2003, p. 572).

There are indeed serious difficulties for the realist to pin down precise ‘content’ of their recipe (to spell out what ‘structure’, for example, amounts to).⁵ It is difficult, first of all, to come up with a well-motivated general recipe in the abstract, without concrete exemplars. Consider the different attempts to characterise structural content in terms of Ramsey-sentences or the ‘semantic view’ of theories, for example, in a way that is detached from actual scientific cases. (See e.g. Cruse and Papineau, 2002; Ladyman, 1998.) It is entirely unclear why these recipes should in general be at all good for discerning something that will be carried over in various theory changes, something that furthermore genuinely accounts for the past theories’ empirical success. (See Saatsi, 2012.) Often recipes are not just defended in the abstract, but along with illustrative exemplars that motivate the recipe and give some concrete content to it (e.g. Worrall (1989) on Fresnel’s ether theory, and Ladyman (2011) on phlogiston). The problem with this use of exemplars, however, is that different exemplars pull in different directions, ultimately watering down the recipe’s content. If by ‘structure’ we mean whatever is common to the Fresnel case, and the phlogiston case, and all the other cases that recipe realist imperialistically wishes to accommodate, then structure as the ‘common denominator’ threatens to become so vague (or disjunctive) that the recipe fails to pin down realist commitments in the way it is purportedly meant to do. All in all, recipe realism can at best be bought for the price of an unacceptable degree of ambiguity in realist commitments.

If one thinks that scientific realism *must* be recipe realism, then so much worse for scientific realism. This is what Stanford’s (2003) ‘trust argu-

⁵An optimist may say that this only shows that the recipe is complex and difficult to figure out. I think it is more realistic to regard the whole programme of recipe realism as a degenerating one.

ment’ concludes, having first demanded that realists absolutely should offer a recipe that could be reliably applied ‘prospectively—in advance of future developments—to identify the idle features or components of scientific theories.’ (p. 915) Stanford argues that we if cannot invent and defend a recipe that allows us to clearly identify trust-worthy aspects of our current theories, then instrumentalism is in the offing. As I will next explain, this is not so; realism without recipes is still realism enough.⁶

4 Realism through exemplars

Recipe realists are right in leaning heavily on exemplars in explicating their realist commitments, but they go wrong in trying to generate a general recipe that captures the gist of those exemplars. In order to bring out the contrast between my realist inclinations and recipe realism it is worth reflecting further on the notion of ‘realism’ itself.

Realism is often construed as a naturalistic philosophical (and sometimes quasi-scientific) *theory*. As a theory about science, realism is broadly speaking explanatory vis-à-vis the empirical success of science, for example, and it is also broadly speaking predictive, by virtue of offering a prospectively applicable algorithm for discerning belief-worthy parts or aspects of our current science. As a unified theory about science it is meant to be testable by data from the historical record and our best understanding of science. (Hence the fairly common talk of counterexamples that ‘falsify’ realism.)

That is how recipe realists construe ‘realism’. I prefer to think of realism as the adoption of a certain positive *attitude towards* science (as opposed to acceptance of a realist *theory about* science).⁷ One’s belief in realism, then, amounts to believing the general appropriateness of this attitude.⁸ The realist attitude is not a theory or recipe, but rather an adherence to the No Miracles intuition according to which the impressive empirical success of science is by and large down to theories latching onto reality in ways that make them empirically successful. The idea that this is a reasonable attitude towards science says nothing in and of itself about where exactly we should place our trust in science; it says nothing about exactly what there is in the world; and so on. But by maintaining this attitude, and acting accordingly in trying to figure out how old theories are related to newer ones, how idealisations work, and so on, one is operating in a realist spirit by virtue of adhering to the quintessentially realist notion that overall

⁶See also Saatsi (forthcoming a) for further discussion of Stanford’s argument in the broader context of ‘pessimistic inductions’ against realism.

⁷My distinction between theory and attitude may have connotations of van Fraassen’s (2002) distinction between factual beliefs and epistemic stances. Examining the relationship between these distinctions calls for further work.

⁸The statement ‘realism is correct’ is a claim not directly about science; it is meta-level claim about the realist attitude, which itself concerns science.

empirical success (and in particular predictive success) of science is a matter of science progressively latching better and better onto unobservable reality.

As a global attitude towards science, scientific realism thus construed has little epistemic content in and of itself. A realist gains her epistemic commitments when she applies her realist attitude locally, in a piecemeal way, to particular theories in particular disciplines and domains of science. Only these further steps will allow her to say something epistemologically committing about the various ways in which we can expect specific theories to latch onto reality. This is where the exemplars come in. They provide cognitive content to claims of the following sort: in a domain of science like *this*, with theories or models like *that*, empirical success in *this* sense, is (probably) accountable in *those* terms (even if these theories or models are radically mistaken ‘on the whole’). In order to fill in the underscored placeholders above a realist can consult the (scientific) experts for the fullness of relevant detail, instead of pretending to be able to figure out the answer from the philosophical armchair. She might focus on the shift from ether theories to Maxwell, for example, and on the successful derivation of Fresnel’s equations more specifically, and fill in the blanks by drawing on our best overall historico-scientific understanding of what made Fresnel’s derivations tick. This exemplar, and many others, can concern historical theory-shifts; others can concern idealised, or internally inconsistent models, for example. The more exemplars she explores, the better grasp she gets of the epistemic commitments that her realist attitude ensues.

By way of analogy, consider our commitment to try to account for biological phenomena from an evolutionary perspective. This commitment, as captured by Theodosius Dobzhansky’s famous phrase ‘Nothing in biology makes sense except in the light of evolution’, expresses an attitude towards our overall understanding of the living world, not a theory about it. Even if one did not have any first-order evidence positively supporting the appropriateness of this attitude, one could take it as a starting point, and defend it in a piecemeal way against anti-evolutionary attacks. The defence varies in detail from case to case: from the accusation that the evolutionary perspective is incompatible with the complexity of flagella, to the claim that it is incompatible with our consciousness, and so on.

Notwithstanding the obvious weaknesses of this analogy, we can similarly start by adopting a broad realist attitude towards science: the impressive success of science, overall, does not make sense except in the light of the idea that science is latching better and better onto unobservable reality. We can then go about defending this attitude against anti-realist attacks. Again, the details of the realist defence will vary from case to case: from the claim that the realist perspective is incompatible with the success of Newtonian physics, to the claim that it is incompatible with the success of Fresnel’s physics, to the charge that it is in tension with our most fundamental theories being mutually incompatible, and so on. Different instances of the realist defence

can rely on different senses of ‘latching onto reality’, and the case-dependent details function as exemplars that help us think about other cases that are relevantly similar to these exemplars.

Let me sum up this brief sketch of exemplar-based realism as follows. In his state-of-the-art review of scientific realism, Chakravartty (2011) begins by summarising scientific realism as ‘a positive epistemic attitude towards the content of our best theories and models, *recommending belief in both observable and unobservable aspects of the world described by the sciences.*’ (My emphasis.) In my view Chakravartty is exactly right in characterising realism as ‘a positive epistemic attitude’, but realism need not make any kind of blanket recommendation regarding what our theories say about the unobservable aspects of world. Instead of this, the positive attitude of exemplar realism recommends only adherence to the idea that by and large theories achieve their empirical success *by virtue of* getting something right about the unobservable world. Specific exemplars then give cognitive content to what this ‘by virtue of’ relation stands for, varying from case to case, with limited projectability to cases relevantly similar to those exemplars.

5 Challenges to exemplar realism?

Realism as a global meta-level attitude does not come with a prospectively applicable recipe for delineating realist commitments. The only forward-looking aspect of realism is associated with the exemplars: an exemplar realist account of the empirical success of an ether model of light, for example, will provide an understanding of how models of *that* kind can be predictively successful in *those* ways (even if the model harbours mistaken assumptions). As is always the case with exemplars, their application is based on (fallible) judgements of relevant similarity. This is a relatively ‘local’ matter. An exemplar from optics is unlikely to be of any use with respect to geology, say, and a realist account of the empirical success of Newtonian gravity, for example, only relates to theorising that is suitably similar to this exemplar, viz. gravity and dynamics in fundamental physics.

This piecemeal spirit of exemplar realism goes against the commonly held idea that realism requires the provision of a prospectively applicable abstract recipe for capturing realist commitments for any given theory. Anti-realists present this as a challenge: if realism does not provide a prospectively applicable recipe, then how do we know what parts or aspects of our current science we can *trust*? (See Stanford, 2003.) In response I maintain that the idea of a prospectively applicable realist recipe is an empty hope. It is a manifestation of philosophical arrogance to think that as a realist philosopher one commits oneself to providing a global recipe—largely independently of the science steeped in relevant details—for revealing what aspects of theory-world correspondence makes any given theory in mature science

tick. A much more reasonable and realistic aim for the realist is to be able to point to a range of exemplars to elucidate the general sense in which the realist hopes to account for empirical success in terms of the incremental progress in theories and models as representations of unobservable reality.

Exemplar realism is at the same time global, in its attitude, and local, in its action. By adopting the realist attitude one does not thereby take on the responsibility to delineate the trust-worthy areas of all of science in one fell swoop. Rather, the realist's attitude is manifested locally. For example, the realist might concern herself with the specific challenge from quantum physics, that despite their most impressive empirical success quantum theories should be construed merely as predictive instruments. Here the realist looks at the relevant details, working with the physicists (or those who work in the relevant 'foundations'), in order to explicate ways in which we can make sense of the empirical success of quantum theories being due to quantum theories latching onto reality. The fruits of that work of course need not apply in any significant way to theories in other areas of science.

It is a necessary consequence of defending realism in this piecemeal way that we must bite the bullet in response to the challenge that with respect to most theories we do not really know what we are realists about. But this is only problematic if one thinks that realism is a theory about the part of the world we call 'science', a theory that tells us what there is in the world, which theoretical terms refer, and so on. Realism as an attitude, as I understand it, is fully compatible with such admission of our epistemic limitations. The piecemeal approach to defending the appropriateness of the attitude gives renewed meaning to the realist slogan *divide et impera*.

At the global meta-level exemplar realism only claims that copious instances of empirical success of science are generally speaking accountable in terms of theories latching onto reality. It leaves it to the exemplars to clarify the different senses in which empirical success can be accountable in realist terms. What does it take to be 'accountable in realist terms'? Clearly this is the critical notion here. One might find this notion unduly ambiguous and vague. In particular, one might challenge the realist (i) to justify her realist reading of the exemplars: by virtue of what do the exemplars furnish *realist* explanations of past successes? And (ii), by virtue of what do these realist readings of the exemplars avoid begging the question against the anti-realist?

The first challenge asks the realist to provide a principled explication of 'a realist explanation' of the empirical success of a theory or model that is clearly false—e.g. a past, now-rejected theory, or a highly idealised model, say. Recipe realists think that this kind of explication can be done via a realist recipe. Arguably the realist is thus pushed towards recipe realism after all, in order to be able to justify her realist reading of the exemplars (the justification being that they fit an independently motivated recipe). In response, the exemplar realist can deny the need to introduce any recipes in

order to get a handle on what it takes to account for the empirical successes of false theories. Regarding past false theories, for example, whether or not these are related to our current best theories in a way that furnishes a realist or anti-realist explanation of a past theory’s empirical success is part and parcel of scientific understanding of those theories. Thoroughly examined cases that are well understood by the lights of science are exemplars that wear their realist reading on their sleeves. Well-understood does not mean ‘simple’, of course: such exemplars can be complex and give rise to very subtle senses in which theoretical correspondence between two theories can be regarded as accounting for the empirical success of one in terms of the other. The relationship between Newtonian gravity to Einstein’s general theory of relativity is a well-studied case in point. Consider the in-depth study of the reasons why, from the point of view of general relativity, Newtonian theory of gravity made various novel predictions and can be used to put a man on the moon. This captures well the way in which one can locally account in a realist spirit for a past theory’s empirical success in extremely rich, *scientifically* kosher terms, in comparison to which realist recipes pale.⁹

The second challenge accuses the realist of begging the question in adopting a realist stance with respect to our current theories in presenting exemplars of realist explanations of past successes. The force of this challenge depends on the overall dialectic. It must be acknowledged that exemplar realism does not even aim to provide a *positive* argument for realism in and of itself. Rather, exemplar realism is a meta-level characterisation of how one can (and should) go about defending global realist attitude against more and less familiar challenges to realism. It is a good question where the prima facie motivation for the realist attitude comes from, but this is not something we have to pack into a statement of exemplar realism. But this is no different from statements of recipe realism, which need not provide a positive argument for realism either. In both cases something has to be said in favour of adopting a realist attitude towards science at the outset, whether it be a general plausibility consideration, or a more specific argument turning on some kind of inference to the best explanation, say. I want to bracket these further issues here.

In closing I will consider one further worry about exemplar realism. Without recipe-like constraints on what realist-explanations of empirical success can look like—the worry goes—realism becomes vacuous. The worry is that without recipe-like constraints it becomes impossible to have counterinstances to realism: for any empirically successful past theory there are *some* reasons why it is empirically successful; pointing to those reasons and calling them an ‘exemplar’ cannot be enough, for otherwise anything goes.

⁹For relevant ‘foundational’ treatment of the relationship between Newtonian gravity and general theory of relativity, see e.g. Malament (2012) and Fletcher (forthcoming). Barrett (2008) and Saatsi (forthcoming b) make the connection to the realism debate.

In response, the exemplar realist agrees that there are constraints for sure, but she insists that these need not take the form of a recipe. The realist can accept (by the same token as above) that there can also be thoroughly examined, well understood exemplars of *non-realist* accounts of empirical success. The understanding of why these accounts are non-realist ones is again part and parcel of proper scientific understanding of the theories and models in question. The accounts can turn on interesting instances of underdetermination or ‘mathematical coincidences’, for example, showing how two different empirically successful scientific theories or models can have little in common beyond their empirical predictions. (See the Kirchhoff case from Saatsi and Vickers (2011), for example.) But such exemplars also have a limited range of projectability, of course, and only speak against the global realist attitude with respect to instances of theorising sufficiently similar to them.

Overall, the global realist attitude regards as relatively rare—but by no means non-existent—such instances of empirical success that ought to be accounted in non-realist terms. The exemplar realist’s piecemeal evaluation of the different cases makes room for a natural sense in which we can be realists, globally speaking, about *most of* science. By virtue of this it is much more robust than recipe-realism.

Acknowledgments A version of this paper was presented at *New Thinking about Scientific Realism* conference in Cape Town. I would like to thank the audience. Special thanks to Steven French.

References

- Belot, G. (2014). Down to earth underdetermination. *Philosophy and Phenomenological Research*.
- Boyd, R. N. (1973). Realism, underdetermination, and a causal theory of evidence. *Noûs* 7(1), 1–12.
- Butterfield, J. (2012). Underdetermination in cosmology: An invitation. *Aristotelian Society Supplementary Volume* 86(1), 1–18.
- Chakravartty, A. (2011). Scientific realism. *Stanford Encyclopedia of Philosophy*.
- Cruse, P. and D. Papineau (2002). Scientific Realism Without Reference. In M. Marsonet (Ed.), *The Problem of Realism*, pp. 174–189. Aldershot: Ashgate.
- French, S. (2014). *The Structure of the World: Metaphysics and Representation*. Oxford University Press.

- Frost-Arnold, G. (2010). The No-Miracles argument for realism: Inference to an unacceptable explanation. *Philosophy of Science* 77(1), 35–58.
- Hofstadter, D. R. (1976). Energy levels and wave functions of bloch electrons in rational and irrational magnetic fields. *Phys. Rev. B* 14, 2239–2249.
- Ladyman, J. (1998). What is Structural Realism? *Studies in History and Philosophy of Science* 29A(3), 409–424.
- Ladyman, J. (2011). Structural realism versus standard scientific realism: The case of phlogiston and dephlogisticated air. *Synthese* (180), 87–101.
- Ladyman, J., D. Ross, D. Spurrett, and J. Collier (2007). *Every thing must go: metaphysics naturalized*. Oxford University Press, USA.
- Peters, D. (2014). What elements of successful scientific theories are the correct targets for ‘selective’ scientific realism? *Philosophy of Science*.
- Putnam, H. (1979). *Mind, Language and Reality: Philosophical Papers, Volume 2*. Cambridge University Press.
- Saatsi, J. (2012). Scientific realism and historical evidence: Shortcomings of the current state of debate. In S. O. S. de Regt, Henk W.; Hartmann (Ed.), *EPSA Philosophy of Science: Amsterdam 2009*, pp. 329–340. Dordrecht: Springer.
- Saatsi, J. (Forthcoming a). Historical inductions: Old and new. *Synthese*.
- Saatsi, J. (Forthcoming b). What is theoretical progress of science?
- Saatsi, J. (Forthcoming c). Explanation and explanationism in science and metaphysics. In M. Slater and Z. Yudell (eds), *Metaphysics and the Philosophy of Science: New Essays*, Oxford University Press, (forthcoming).
- Saatsi, J. and P. Vickers (2011). Miraculous success? Inconsistency and untruth in Kirchhoff’s diffraction theory. *The British Journal for the Philosophy of Science* 62(1), 29–46.
- Smeenk, C. (2013). Philosophy of cosmology. In R. Batterman (Ed.), *The Oxford handbook of philosophy of physics*, pp. 607–652. Oxford: Oxford University Press.
- Stanford, P. K. (2003). Pyrrhic Victories for Scientific Realism. *Journal of Philosophy* 100(11), 553–572.
- van Fraassen, B. (2002). *The Empirical Stance*. London & New Haven: Yale University Press.
- Vickers, P. (2013). A confrontation of convergent realism. *Philosophy of Science* 80(2), pp. 189–211.

Worrall, J. (1989). Structural Realism: The Best Of Both Worlds? *Dialectica* 43, 99–124. Reprinted in Papineau, D. (ed.) *The Philosophy of Science* (Oxford: Oxford University Press), pp. 139–165.