A

REALISM, ANTI-REALISM, AND THE SUCCESS OF SCIENCE

by

David Shein

A dissertation submitted to the Graduate Faculty in Philosophy in partial fulfillment of the requirements for the degree of Doctor of Philosophy, The City University of New York

2002

UMI Number: 3037445

ProQuest Information and Learning Company
300 North Zeeb Road
P.O. Box 1346
Ann Arbor, MI 48106-1346

This manuscript has been read and accepted for the Graduate Faculty in Philosophy in satisfaction of the dissertation requirement for the degree of Doctor of Philosophy.

1/31/02
date

Jonathan Adler
Chair of Examining Committee

1/31/02
date

Michael Devitt
Executive Officer

John Greenwood

Arnold Koslow

Michael Levin

Supervisory Committee

THE CITY UNIVERSITY OF NEW YORK

Abstract

REALISM, ANTI-REALISM, AND THE SUCCESS OF SCIENCE

by

David Shein

Adviser: Arnold Koslow

According to the success argument for scientific realism, scientific realism must be true or else the success of science would be inexplicable. I contend that this argument fails: scientific realism is unable to account for the predictive successes of scientific theories. Scientific anti-realism also fails to explain these successes, as does social constructivism. I conclude that we are forced to take the second horn of the realist's dilemma and accept that the success of science is inexplicable.

ACKNOWLEDGEMENTS

Somewhere between New York City and Annandale, between a 5th floor walk-up in Hell's Kitchen and a house in the country, between Rudy's and a daughter, I managed to finish my thesis. Lots of friends along the way helped.

Arnold Koslow, Michael Levin, and John Greenwood helped with the writing and thinking and re-writing and re-thinking that gave shape to the finished product (though neither the errors committed nor the positions endorsed ought to be attributed to them). They helped me to become a better thinker, writer, and philosopher, and I thank them for that.

I was lucky to forge close friendships in graduate school with Martin Harvey and Daniel Kaufman, friends who have listened when I needed to talk and talked when I needed to listen; conversations with them over the years helped me figure out what I was trying to say in this project. I talked about the nature of realism with Bradley Armour-Garb, another friend from graduate school, in person, over e-mail, on the telephone, and on the racquetball court; these conversations helped solidify my thinking about the early parts of the thesis. It gladdens me that these are still friends and interlocutors.

Rebecca Holsen was my junior year adviser when I was an undergraduate at SUNY Oswego; she told me I was bright but lazy and pushed me to work harder. Then she went to California. Paul Blunt was my senior year adviser; he taught me *how* to work harder and how to *do* philosophy. Without Rebecca and Paul, I would never have made it to graduate school.

Miriam Lahey, friend and mentor from Lehman College, reminded me, time and again, that I was still in graduate school, and she made it possible for me to attend to my thesis while working full-time. Jonathan Becker did much the same thing for me when I arrived in Annandale. I was lucky to find one job where my boss became my friend; how extraordinary to find another.

Thank you to Vibe Pape, who helped compile the bibliography and caught some potentially embarassing grammatical errors. Thanks also to Christopher Altman for photocopying and collating.

My mother and father never quite understood why I wanted to study philosophy but, to their credit, they supported me anyway. So, too, for my brother and sister. Mom, Dad: thank you all for your support. I'm finally done with school.

My dear wife, Amy, did understand why I wanted to study philosophy and supported me anyway. *That's* love. Amy, thank you for putting up with my 5 a.m. alarm clocks, surly attitudes, paranoia, frustration, moodiness, and distraction. Time and again, you have been better to me than anyone deserves. The thesis is done; at last, there are just the three of us in the house.

Meyer Hines and David Falkenstein: I wish you could have read this.

Emily Rose, this is for you. When you're old enough, maybe you'll read it and understand why I sometimes had to stop playing with you and sit at the computer. Now we can play all day.

TABLE OF CONTENTS

INTRODUCTION

Scientific realism is one of those theories that seems just obviously true. What could be more evident than the fact that science tells us what the world is like? With a staggeringly small number of exceptions, everyone is a scientific realist. To make matters worse, the opponents of realism are more often than not the same crackpots who believe in alien abduction theory and parapsychology. Given the sheer intuitive strength of the position, scientific realism is hardly in need of more support. What it needs is a critical appraisal. The present project seeks to offer that appraisal.

As Putnam (1975) points out, there are two types of arguments for scientific realism. The negative argument is the failure of existing scientific anti-realist programs. Putnam lists instrumentalism and operationalism; we can add to that list idealism and social constructivism. The positive argument, Putnam notes, is the success argument: the only viable explanation we have of the success of science is scientific realism. So scientific realism must be true. It is this latter argument which is the subject of the present essay: I want to critically assess this argument for scientific realism. I will contend that, on the most reasonable construal, realism fails to offer a cogent account of the success of science. I will then argue that neither anti-realism nor irrealism can account for success, either; the success of science, at the end of the philosophic day, is inexplicable.

Our point of departure is Putnam's declaration that scientific realism is the only philosophy of science that doesn't make the success of science a miracle. This is, of course, a variant of Smart's cosmic coincidence argument for scientific realism. Putnam and Smart are not alone in their adoption of this argument. Michael Levin (1984: 124) notes that "Many philosophers call themselves scientific realists because realism appears to offer the only explanation for the predictive success of science.... [R]ealism predicts that 'true' theories will work; 'true' theories do work; hence realism- or so it seems." All of the articles in Leplin's seminal (1984) are concerned with "what [scientific] success consists in, how it is to be explained, and the role of realism in explanation" (1984: 1), and Leplin's (1997) is concerned entirely with a version of the argument. Peter Lipton (1991: 158-9) notes that "there is a well known argument for scientific realism ... the argument is that we ought to infer that scientific theories that are predictively successful are (approximately) true, since their truth would be the best explanation of their success." Clifford Hooker (1987: 322) points out that "recently, a variety of [...] explanatory theses have been urged on realism, namely that it should explain the explanatory/predictive success of science and the success of realist methods in science." Robert Klee (1999: 313) calls the success argument "the single most powerful argument ever concocted on behalf of realism." Michael Devitt (1991: 108) takes it to be the only argument there is for scientific realism.

Clearly, this is a popular argumentative strategy: 'scientific realism must be true because only scientific realism can explain the success of science.' The argument gains much of its strength from the assumption that scientific anti-realism is unable to explain the success of science- that Putnam's negative argument for realism goes through. On the face of things, this is a reasonable assumption to make. The main anti-realist programs currently being defended in the literature are van Fraassen's constructive empiricism and

Fine's NOA. Their explanations of the success of science are inadequate, though. van Fraassen argues that success can be explained in evolutionary terms: the reason current theories are successful is that only successful theories survive. But this fails to define any feature of our theories in virtue of which they are successful; we are left wondering why these theories rather than some other ones have survived. Fine argues that success can be explained by a non-metaphysical appeal to the accuracy of science's representation of the world – science is successful because it gets the world right – but that this claim should not be taken to be an endorsement of realism. To most, this sounds like a tepid form of realism. Add to these failures the assumption that instrumentalism and operationalism are failed programs, and add the near-universal disdain for constructivism, and you wind up with Putnam's negative argument: there are no viable anti-realist options when it comes to explaining the success of science.

The intuitive force of the success argument is, to be sure, overwhelming. But intuitions are not arguments, and there are grounds for criticizing both the positive and the negative arguments for realism. The first two chapters of the essay offer a prelude to this critique: the first chapter is devoted to a discussion of the several conceptions of realism and anti-realism, and the second one offers a discussion of the several conceptions of success. I then turn my attention to the failures of van Fraassen's and Fine's accounts of success (chapter three) and of the realist account of success (chapter four). The conclusion of chapter three suggests that Putnam's negative argument is right: existing anti-realisms cannot account for success. But the conclusion of chapter four suggests that realism cannot account for success, either; the positive argument for realism fails to go through.

Having cleared the way for alternative accounts of success, I proceed in the fifth chapter to examine sociological approaches to the problem of the success of science. In the final chapter, I sketch a version of constructivism that seems most likely to account for success but conclude that it, too, fails to do the job. In the end, neither philosophy nor sociology appears able to explain the success of science.

The examination of constructivism in chapter six offers an interesting meta-philosophical note on which the essay ends. The problem with realism, I argue in chapter four, is that is cannot explain why false theories succeed (or why false theories apparently succeed); I argue in chapter six that constructivism cannot explain why true theories succeed. There arises a peculiar symmetry of explanatory failure: sociology accounts for false theories' successes but not true theories' successes, and philosophy accounts for true theories' successes but not false theories' successes. This has implications (noted in chapter six) for the symmetry principle upon which the sociology of science is based. It also has implications, I think – though I'm not sure what they are – for the philosophy of science. Thus the end of this work on scientific realism is the beginning of work to be done on the philosophy of science.

It's worth mentioning that this essay is not about either the rationalism / irrationalism debate or the thesis of under-determination. Both of those are topics about which much has been written, and both are tangential to the realism issue.[1] The rationalism / non-

[1] John Greenwood points out (correctly) that this claim (like the one that follows it) makes little sense in the absence of an account of realism. Fair enough. I ask the reader's indulgence, though: an account of realism will be provided in the next chapter. For now, we can rest with the rough formulation given in the text.

rationalism debate is concerned with the question of whether theory change can be explained in purely rational terms. This is not the realism / anti-realism debate: realism is concerned with the relation between science and the world, while rationalism is concerned with the relations between theories (and the reasons for changing theories). Nevertheless, there are obvious points of overlap: one very good reason for changing theories is the justified belief that T_2, but not T_1, accurately represents the world. But there is no obvious argumentative route from rationalism to realism: rationalism presupposes realism, so one cannot argue from rationalism to realism.[2] Given our purposes of assessing the success argument for realism, then, it will do us no good to argue for or against rationalism.

What about the thesis of under-determination? To believe in the under-determination thesis is to believe that all the available evidence is compatible with multiple, incompatible theories. This might provide grounds for being a scientific anti-realist: if you think that, for any theory T there is an incompatible theory T_2 that can account for all for which T accounts, then you are likely to despair of ever finding out which theories get the world right and which ones do not. This is, I think, a legitimate way of arguing for anti-realism. What's more, there is a link between it and the success argument: one could run the two together in favor of anti-realism ('success fails to prove realism because, for any theory can be shown to be successful, an alternative can be offered) or in favor of realism (for some successful theories, there is no substantive alternative, and this is proof that anti-realism is wrong'). But – and this is the point – one can endorse or reject the

[2] Rationalism, so understood, anyway. One might have rational grounds for preferring one theory to another without being a realist. In this case, rationalism and realism are not even tangentially related.

success argument for reasons having nothing to do with the thesis of under-determination. Our analysis of anti-realism and the success argument need not involve a defense of the under-determination thesis.

In the spirit of separating out issues, we should note the connection between the debate over realism and the debate over nominalism. There is, again, an overlap: if you think objects have their properties essentially and so independently of observation, you are likely to believe that those objects exist independently of observation. Again, though, the overlap is not compelling: one could think that entities exist independently of human perception but that their putative properties are the result of perception. This means that the issue of inherency, like those of rationalism and empirical equivalence, is tangential to our present study and so need not occupy a central place in it.[3]

It remains to be said why this inquiry into the viability of scientific anti-realism should be of interest to anyone. If all parties to the debate agree that science is successful, why should we care to explain that success? A brief look at the history of science teaches us that science's successes have been impressive, but it teaches us also that it's failures have been equally impressive. These failures are not necessarily sufficient reason to give up realism, but they should instill in the critical observer a certain measure of skepticism. The debate over scientific realism begins when we begin the self-reflective work of deciding whether this skepticism is warranted. If we determine that it is not, we need to show how to dispel the skepticism. If we determine that it is, we need to do the anti-

(Thanks to John Greenwood for this suggestion.)

realist work of accounting for success in other terms. All of this can be done within the context of accepting success. Science is a tool we have for measuring, explaining, and manipulating the world around us. The philosophy of science seeks to understand how that tool works, and this can be done without doubting that it works.

It will be objected in some quarters that the inquiry can be forestalled by pointing to the realism of working scientists. If we are seeking only to explain how science works -- if we are not trying to change science or make recommendations to its practitioners -- then we can solve our problem by saying that we should be no more or less skeptical than scientists are, and we should handle that skepticism just as they do. This is to fundamentally mis-understand the goal of the philosophy of science. Like science itself, it seeks not to recommend, but to explain.[3] The scientist explains why the ball in the wagon looks like it rolls to the back when the boy starts pulling on it. He doesn't recommend anything- he just explains why the world appears to us this way. So, too, for the philosopher of science: he doesn't want to change anything; he just wants to explain why science appears to be so successful. If, like the ball in the wagon, things aren't exactly as they appear, so be it. *That's* the value of the inquiry.

---

[3] For a discussion of these and related issues, see Wartofsky (1991), who takes the key issue for realism to be the question of how to be a realist without being an essentialist.

1. REALISM

The realist argues that realism must be true because realism offers the only viable explanation of the success of science. This essay is an extended critical analysis of this claim. The present chapter seeks to determine what belief or set of beliefs the epithet 'realism' is supposed to pick out; this is a necessary first step in dissecting the claim that only it can explain success.

There are three broad conceptions of realism: the semantic conception, the epistemic conception, and the ontological conception. Contemporary discussions favor the first two conceptions, with the ontological conception regarded as naïve and outdated. It is the thesis of this chapter that, contra popular opinion, the ontological conception of realism is the one which is best able to get the realist's explanation of success off the ground. The semantic conception is parasitic upon the ontological one, and the epistemic conception cannot play the explanatory role required of it by the success argument. I defend the ontological conception against charges of naïveté, contending that scientific realism gets its realist bite from its ontological commitments -- the semantic and epistemic commitments that accompany these ontological ones merely facilitate discussion -- and the success argument for realism makes the most sense when these ontological underpinnings are brought to the surface.

---

[4] Barnes, Bloor, and Henry (1996:30) remark, on the related issue of the connection between science and the sociology of science that "For the scientist the world is the object of study; for the sociologist it is the

1.1 the semantic conception of realism

Contemporary discussions of scientific realism tend to take that position to be a semantic thesis- something like 'a realist takes the central terms of mature scientific theories of science to be genuinely referring expressions', or 'a realist takes the sentences of the mature theories of science to be true'.[5] A not-so-popular conception of realism takes it to be a metaphysical or ontological thesis: 'a realist thinks that (some or most of) the entities postulated by the mature theories of science really exist.'[6]

The argument for rejecting the ontological conception in favor of the semantic one is quick: we cannot make sense of this 'really exists' locution; the best we can do is talk about the truth conditions for sentences about the entities. So a better formulation of the realist position will center not on the existence of the entities, but on the truth conditions for sentences about the entities. Specifically, a realist is someone who thinks that the truth conditions for such sentences are independent of speakers of the sentence- we recast 'really exists' as 'has independent truth conditions.' Realism thus becomes a commitment to the correspondence theory of truth, and the debate over scientific realism transforms into a debate over the nature of truth.

---

scientist-studying-the-world that is the object."

[5] John Greenwood points out that there is a weaker semantic conception of realism available, *viz.*, the view that theories purport to refer. But this view cannot underwrite the success argument: that a theory merely purports to pick out things in the world cannot explain why the theory issues in successful predictions and explanations.

[6] Canonical statements of the referential conception are found in Boyd (1973, 1984) and Putnam (1975, 1984). The conception in terms of truth is articulated by Laudan (1984), Dummett (1993), by Boyd and by Putnam. The ontological conception of realism is mentioned favorably by Hacking (1983) and defended by Levin (1984) and Devitt (1991).

I think it's a mistake for the realist to make this move. It's possible that the proper analysis of truth will turn out to be something other than the correspondence theory. But the serious realist (about anything except truth) would not change his commitments or rethink his faith in science- he would not stop believing in the existence of those things he used to think exist, and he would not stop regarding as effective those methods of inquiry that he used to regard as effective. What theory of truth one holds is irrelevant to one's commitment to realism. Indeed, we have only to look to the history of philosophy to find realists who are not committed to any theory of truth-- realism predates the debate over the nature of truth! Of course, you can argue that these realists had an implicit commitment to a theory of truth, but this is to give away too much: it is to admit that the semantic conception of realism is merely a variation on the ontological conception, with the move to truth conditions playing no greater a role than that of semantic ascent. Of course, semantic ascent is okay, but it's metaphysically toothless: a realist who wants to talk in a reasonably sophisticated manner about some class of entities will revert to talk of truth conditions for statements about those entities, but what makes this person a realist is the commitment to the entities, not the commitment to the manner of talking about the entities.

This analysis presuppose a particular, metaphysical conception of realism. Anyone who rejects that conception is likely to reject the arguments of the previous paragraph as question-begging. So let me key my rejection of the semantic conception of realism to the success argument. The realist claims that realism is the only explanation of the success of science. The suggestion that realism is a semantic doctrine requires us to commit to something like the following: what explains the success of science is the fact that the statements of our scientific theories are true. But it should be immediately obvious that any explanatory potency is gained not from the mere truth of the sentences, but from the fact that the world is the way the sentences say it is. It's not the nature of the truth conditions that explains success; it's the fact that the conditions obtain.[7] What makes an airplane stay up, for instance, is not the truth of our theory of aerodynamics; it is the fact that "the pressure on the underside of a moving airfoil is greater than the pressure on its overside" (Levin 126). It's not the truth of a theory that explains its success; it's the theory itself. The theory offers an account of the goings-on of some part of the world, the world works in that manner, and so the theory succeeds. No truth needed.

But,' the semanticist points out, 'this talk of "accounting for the goings-on of some part of the world" is semantic - "accounting for" is just a variant on "represents", and that is unabashedly semantic.' 'Indeed,' she might argue, 'to say that a theory is true is just to say that it accurately represents the world.' The intended conclusion, of course, is that my account of realism *is* semantic.[8] This analysis is both correct and misleading. Imagine the dialectic: the realist says that our scientific theories are successful, and we ask why. The response is that that they are true. We ask what this means, and we are told that the theories accurately represent the world. This provides us with both an account of the nature of truth and an explanation of the success of the theories. It's not merely the fact that the theories are true that explains their success; it's the fact that the theories are true

[7] As Michael Levin says, "truth, like Mae West's goodness, has nothing to do with it" (1984: 126).

and that this is what it means to be true. Insofar as this is the case, the semanticist explanation of success is correct. But now imagine that someone else comes along and takes issue with the analysis of truth. A debate ensues over the what it means to say something is true, and the realist ends up defending the correspondence theory of truth. The terms of the debate thus shift: we are no longer interested in explaining the success of our theories ('they accurately represent the world'); now we are concerned with the question of whether to call this property of theories 'truth.' It is in this sense that the semantic reading is misleading. What the realist should say here is that she stands by her explanation of the success of science -- call it what you like -- and she also stands by her analysis of truth, and she will defend herself on both fronts. What's lost in the semantic analysis of realism is the fact that these are separate fronts: the theory of truth and the explanation of success are two separate theoretical commitments.[9] My rejection of the semantic construal of realism is motivated primarily by a fear of this sort of slide between issues: let those who want to discuss the nature of truth fight about it on their own; the realist explains success by appealing to the fact that successful theories accurately represent the world, regardless of whether this earns those theories the badge of truth.[10]

---

[8] I'm not making this up. Brad Armour-Garb (a realist and a semanticist) recently ran pretty much a verbatim version of this argument on me

[9] Brian Ellis (1985) finds three key components in the thesis of realism: a commitment to a physicalist ontology, the thesis that "the laws and theories of science are genuine claims about reality" ('the central thesis of scientific realism'), and "the objectivity thesis that the laws and theories of science are objectively true or false" (52-53). He remarks that "the objectivity thesis is rarely distinguished from the central thesis of scientific realism and is often confused with it. But [...] they are quite distinct" (53). And he goes on to argue that the usual realist commitment to the correspondence theory of truth lies at the heart of the confusion and that what makes someone a realist is not a commitment to this theory of truth but, rather, a commitment to "the idea that there are things in the world to which our laws and theories refer and of which they are true or false" (52).

[10] This isn't to say that all anti-semanticism is motivated by such fears. van Fraassen rejects appeals to truth because of his epistemology. Fine, Levin, and Hacking think that theories are themselves explanatory and that appeals to the truth of the theories do not increase their explanatoriness. Redundancy theorists treat truth simply as a non-explanatory property. The argument in the text grants the explanatoriness of truth

Michael Dummett (1982, 1993) treats the realism dispute as a dispute over truth conditions. If you think this then you are going to reject the reason given above for not being a semanticist, *viz.*, that one's metaphysical commitments will remain the same regardless of the analysis of truth. This might be the case, Dummett would say, but it's irrelevant to realism. If what one thinks is the method of determining the truth of a scientific claim is what determines one's position on realism, then the analysis of truth *is* relevant to realism, and metaphysical commitments are *not*, and Dummett's right, and I'm wrong.

I won't attempt to refute Dummett's analysis of what it means to be a realist; it would take us too far afield. I think it's fair to say that his read on the nature of realism is idiosyncratic (while popular), and not sensitive to the history of realism. But we can let that slide. Suffice it to say that the Dummettian realist is unable to run the success argument for realism. What explains the success of science? On the Dummettian view, the realist would argue that what explains the success of science is the fact that bivalence holds for the sentences of science. Well, that just doesn't make sense. And it does no good to point out that Dummett is an anti-realist. The point is that his account of realism literally makes no sense when considered in the context of the argument from success. It's not that realism, taken in these terms, is explanatorily impotent. Rather, it's that we can't figure out how bivalence could *be* explanatorily potent or impotent. Maybe Dummett's right about the connection between truth and realism; if he is, the success

---

and argues against semanticism on other grounds; one could argue against it for any of these reasons as well.

argument drops out. Insofar as we are concerned in this work with the success argument, then, it can't be Dummett's conception of realism with which we are concerned.

What about the version of the semantic conception of realism that focuses on reference rather than truth? The reason reference is held to matter is that science is believed to be convergent: scientific theories build on one another, improving on past theories in the process. What stands in need of explanation is the convergence of scientific theories on the truth, and what is supposed to explain this is the preservation of reference across theories. *Vis-à-vis* the success argument, this works out as follows: the success of science is the convergence of scientific theories, and what does the explaining is the preservation of reference. This makes scientific realism the thesis that the central terms of scientific theories genuinely refer.[11]

The problem with this argument is that it seems to get wrong both the *explanans* and the *explanadum*. When we talk about the success of science standing in need of explanation, we don't have in mind anything like convergence. Rather, we have in mind things like explanation, prediction and manipulation of our environment. (Indeed, Putnam (1975: 141) says that this is what he is talking about when he talks about success.) And reference is not able to account for these types of success. I've got this theory: the reason the computer won't work is that the monitor is not connected to the hard drive. You've got this other theory: the reason the computer won't work is that the fuse has blown in the

[11] Putnam introduces a principle of charity to make the preservation of reference work. I'll have a little to say about this in a couple of paragraphs, and a lot to say about it in chapter four. For present purposes, I propose to give the realist whatever tools he needs to make reference work: my point is simply that reference and realism are different creatures, regardless of the details of the theory of reference.

office. What makes your theory successful and mine unsuccessful is the fact that your theory gets things right and mine does not. [12] The fuse really is blown, and this is why the computer won't turn on. Reference has got nothing to do with it, except in the utterly trivial sense that our terms pick out the same things. What explains the success of a theory is the fact that the theory says that x, and x is the case. Of course, the theory's terms have got to connect to the right things in the world, but this connection isn't what explains the success- the fact that the world is the way the theories say it is what explains success.

But what about the convergence of theories? What explains why one theory is better than (but still resembles) the theory that preceded it? Part of the answer to this question is that the new theory takes the central parts of the older one and refines them. This process of refinement requires that the referential relations of the old theory be carried over into the new theory; reference must be preserved to explain the convergence of the theories. This explains why the new theory is to be considered an improvement on, and not a replacement of, the old theory. But what explains why it's an improvement (rather than a step backwards, or no better than the old theory) is the fact that the new theory gets right certain facts about the world that the old theory didn't. Preservation of reference explains why we have convergence of theories rather than divergence, but preservation of reference fails to explain why we have convergence on truth. Alternatively: we are able to call the new theory an improvement on the old theory because there is a preservation

[12] I know they're not theories in any global sense of the term. But they serve to make the point, and the point is extendable to any theory you like. And don't get conned into thinking about this in terms of theory choice: we can test the theories, run counterfactual analyses, etc. That's all irrelevant to the issue of what explains success.

of reference -- this is what makes them theories about the same thing -- but to regard the preservation of reference as the explanation of the improvement is confused.

Underlying this confusion is, I think, a confusion over two types of convergence. When we say that two theories are convergent, we might mean that one of the theories captures and extends the content of the other theory- that the theories are points on a single trajectory (rather than points on distinct trajectories). For this to be the case, there must be preservation of reference, for this is what allows the newer theory to extend the older one. When we say that two theories converge on the truth, however, we mean that the theories both get right some fact about the world. The first type of convergence is a convergence of theories with one another; the second type is a convergence of theories on the truth. The success argument requires the second type of convergence, but preservation of reference is required only for the first type.

In addition to this worry about the semantic (referential) conception of realism, there is a familiar, more localized problem. As Laudan (1984) points out, the historical record contains lots of theories that have been successful but were non-referring. This is taken to show that reference is not required to explain success, and Laudan concludes that he has countered the success argument for realism.[13] Much has been written on Laudan's pessimistic induction, and I will have something to say about it in my chapter four discussion of realism. For now, though, it is worth looking at the threat posed to the semantic conception by Laudan's argument. His conclusion is that the success of science does not require genuine reference. At first glance, Putnam's referential gambit seems immune to this criticism. We should construe the referring expressions of past theories in such a way that they come out as cases of successful reference- we should take them to refer to whatever current theories say exists and play the proper roles in the corresponding theories. Of course, theories that were wholly false (like phlogiston theory) are not going to be held to refer. One wants to run him onto a slippery slope here, contending that there is no reason to withhold charity from phlogiston but to grant it to other theories, thereby pushing the conclusion that all theories are genuinely referring expressions and so realism is true about everything. This surely would constitute a *reductio* on the referential conception of realism (or on the theory of reference).

While Putnam is silent on this matter, we can help him out by appealing to the distinction Kitcher (1993: 140-149) draws between working and idle parts of theories and suggesting that what gets preserved across theories are the central terms of the theories- the theories' working parts.[14] An empirical claim is then offered: the working parts of past, false theories all contained genuinely referring working terms, and those terms are preserved across theories. So what explains the convergence of theories is the fact that their working parts contain genuinely referring expressions, and that reference is passed onto to successor theories.

---

[13] Frustratingly, Laudan claims to have countered an argument for *epistemic* realism. On my accounting, he has countered an argument for *semantic* realism, since he has argued against a view that takes genuine reference to be characteristic of realism. I'll talk more about this in the next section on epistemic realism.
[14] Leplin (1997: 69) draws a similar distinction between essential and non-essential propositions. In developing an account of novelty, he argues that, in order for a result to be novel for a theory, it must be non-essentially involved in the rational re-construction of the development of the theory. This isn't exactly

I'm suspicious of this move. How do we determine the working parts of an historical theory? We could appeal to the explicit statements of the theory, as those theories are offered in the historical record. This is what Kitcher (1993: 145-149) does. But this is overly restrictive: certainly we don't want to say that what makes a theory work is all and only those things that advocates of the theory say make it work. This is the antithesis of Putnam's principle of charity! The only other option, though, is to determine what counts as a working part on the basis of what current theory dictates. But this just iterates the principle of charity: we are asked to re-construe the reference of past theories on the basis of current theory *and* we are asked to re-construe the essential and non-essential parts of those theories in the same manner. Do past theories really refer? Sure; we just have to re-cast them in light of current theory. Ah, but what about theories that were wholly non-referential? They referred too; we just have to re-interpret them in light of current theory. No one who is unimpressed with Putnam's defense will be impressed with Kitcher's save; it's the same move.

I'll defend Laudan's attack on realism later on. At this stage, the point I want to make is not that realism is wrong, but that it is wrong to conceive of realism in semantic terms. Laudan's pessimistic induction shows that the realist who casts her realism in semantic terms cannot explain success. What's more, I have argued, the argument from success for this sort of realism turns on an equivocation. Finally, there is no motivation to go semantic: there is nothing in the success of science that requires us to construe realism in semantic terms. What explains the success of our theories is the fact that the world is the

---

the same as Kitcher's distinction, but it smacks of the same sort of Whiggism. I'll have more to say about this later.

way they say it is. The theory says that there will be an eclipse, there is an eclipse, the theory is successful. This is semantic only in the trivial sense that everything we do is semantic; what explains the success of the theory (say, compared to the unsuccessful cosmology of another theory) is the fact that it gets the world right (and the other theory doesn't). I know, I know: our theories contain genuinely referring terms (and theirs don't). But this is just so much cleansing of ontological hands. To say that the terms are genuinely referring is just to say that there exists what the theories say exists. If you want to do this in terms of reference, okay. But don't tell me that the commitment to a theory of reference is what makes you a realist; what makes you a realist is the fact that you think there is something to which your terms refer.[15]

1.2 the epistemic conception of realism

In the previous section, I spent a fair amount of time analyzing the semantic conception of realism, considering and rejecting several strains of the argument. In this section, I shall be more brief. The dialectical reason for this is simple: more people subscribe to the semantic view of realism than do the epistemic view of realism. Nevertheless, there is at least one vocal realist who defends epistemic realism, and there is at least one vocal anti-realist who claims to have shown epistemic realism wrong.[16] So it's worth looking at and saying why I think it's wrong. My strategy will be to argue that the success argument makes no sense when realism is construed as an epistemic thesis.

---

[15] Not a few semanticists are ready to concede this point. Yet they insist that reference is what makes realism work. The situation is akin to the one described in our earlier discussion of truth: there is a theory of reference that is simpatico with realism, and the realist confuses the two. It's not the commitment to the

Epistemic realists typically contend that it is reasonable to believe, or that we are justified in believing the statements of science. This is, like its semantic cousin, contrasted with the ontological conception of realism, according to which the entities and processes postulated by science exist. Geoffrey Hellman (1982: 245) offers a statement of the epistemic conception, casting realism as the position that it is reasonable to believe at least some of the theoretical statements of science. Jarett Leplin (Leplin 1997: 26) defends a minimal epistemic realism (MER), describing it as the belief that "there are possible empirical conditions, realizable in principle, under which we would be justified in judging some deep structural statements to be true" Later on, he offers a more straightforward statement of this view: "...there are both systematic features of scientific method and specific programs of successful research that make no sense unless beliefs in the reality of theoretical entities, and the possibility of determining their properties empirically, are attributed to researchers" (169).

The first thing to notice is the epistemic conception of realism corrects the motivational error of the semantic conception: it does not seek to sanitize our discourse about what exists. Instead, it shifts the topic of discussion from existence to belief. Indeed, it is a reasonably short step from the epistemic conception to the ontological one- all that is needed is a principle of ontological commitment that moves from reasons to believe in the existence of a thing to a commitment to the existence of that thing.[17] The epistemic

theory of reference that makes the realist a realist, even though the theory of reference complements the realism.

[16] The references are to Leplin (1997) and Laudan (1984), respectively.

[17] To be sure, Leplin notes that his statement of the epistemic realist position is a statement of a minimal epistemic realism, and he acknowledges that "[m]ore ambitious epistemic realisms go so far as to endorse some of the theoretical structures of extant science" (26). I would paint him as an ontological realist, then, who gets there via epistemology rather than abduction. This doesn't matter a whole lot, but I think it points

realist is hardly running scared from metaphysics; she just thinks that the epistemic claim is all that is needed to make one a realist.

This matters. Someone is an epistemic realist just in case he thinks belief in the existence of the theoretical entities of science is warranted. This is consistent with the total falsity of the theoretical picture of science: since false beliefs can be warranted, the scientist can have reason to believe in the existence of theoretical entities that do not exist. This does capture nicely the epistemic features of the position: it is neither semantic (though of course we could construct a semantico-epistemic version of the position) nor ontological (in the sense of requiring that the entities actually exist). Our problem with the semantic version, recall, was that it required the ontological commitments of realism to do all the work, while at the same time trying to keep ontology at bay. Epistemic realism, by contrast, has no truck with ontology whatsoever; the onus of realism is instead on justified belief.

The second thing to notice -- and this is an extension of the first thing -- is that the epistemic conception is weaker than the ontological one. The epistemic realist is committed only to reasons to believe the entities exist, while the ontological realist is committed to the actual existence of the things. And here's my argument against the epistemic conception: This weaker conception of realism is too weak to underwrite the success argument for realism. It might be a perfectly good conception of realism, but it can't be what the realist means when he says that scientific realism, but not scientific anti-

up a parasitism of the epistemic conception on the ontological one. This will come up again in a few paragraphs.

realism, can explain success. The argument from success has it that realism offers the only explanation of the success of science. If realism is taken as an epistemic thesis, then the argument from success would have it that the only explanation of the success of science is the fact that we are justified in believing the entities of science exist. But this can't be right; presumably, science could be successful even if we weren't justified in believing in the existence of its entities. Indeed, realists often point to the novel success of scientific theories as evidence of scientific realism. But we haven't any justification for believing in the existence of something heretofore not known to exist. So justified belief cannot explain success, and so the proponent of the success argument cannot have in mind the epistemic conception of realism.

I do not intend this to be an argument against scientific realism. It would count as such only on the assumption that the success argument is the only argument there is for realism. I am not making that assumption. The focus of this essay is much narrower: I am seeking to dissect a particular argument for realism. And the focus of this chapter is even more narrow: I am seeking to figure out what conception of realism is at play in the argument I am seeking to dissect. The present point is that, on pain of incoherence, it can't be the epistemic conception. This leaves plenty of room for the realist to formulate a different argument for realism, and I think Leplin is probably on to something in his move from reasonableness of belief to ontological commitment.[18] Of course, what counts

[18] van Fraassen is onto the same thing: this is the point of his well-known claim that it's not an epistemic principle to hang for a sheep as well as a lamb (1980: 72). He thinks that if you make your (scientific) ontological commitments solely on the basis of reasonable belief, you will wind up a constructive empiricist. Leplin (on my reconstruction) thinks that this approach leads to realism. John Greenwood, too, notes that if you grant epistemic realism sufficient strength, realism doesn't need the success argument. Here we have the basis for an argument over realism that doesn't turn on the success of science. Like I said, the success argument isn't the only argument for realism.

as a reasonable belief is a can of worms unto itself but -- and this is my point -- it's not a can that needs to be opened in the present context.

It would be appropriate at this point to engage Leplin's (1997) argument for epistemic realism. Since I am going to argue against scientific realism in a later chapter, I will limit my remarks here to a defense of the claim that epistemic realism is too weak to underwrite the success argument. Leplin argues that the type of success at issue in the success argument is novel predictive success, and he argues that a prediction is novel for a theory just in case the prediction is independent of the theory (that is, is not essentially involved in a rational reconstruction of the reasoning that produced the theory) and unique to the theory (that is, is not predicted by any other theories on offer at the time).[19] He concludes that epistemic realism is alone up to the task of explaining such success. The argument turns on the claim that the only alternative to attributing at least partial truth to predictively successful theories is acquiescence in mystery; in other words, Leplin's argument for minimal epistemic realism is the success argument.[20]

Notice, however, that the argument actually employed by Leplin is stronger than he lets on: the alternative to mystery is the attribution of truth to theories, while MER requires only that such attributions be warranted. What explains success, then, is not the mere warrant to interpret theories realistically, but actually doing so. Epistemic realism has collapsed into semantic realism and, as we have seen, semantic realism gets its explanatory power from its ontological underpinnings. In short, epistemic realism is

[19] Leplin's analysis is far more detailed than this. See chapter three of his (1997) for a complete statement of his account.

unable to explain novel success. What does the explaining is a surreptitious commitment to ontological realism.

Leplin seems to realize as much. He writes: "...MER itself is not proffered as an explanation of anything. Rather, realism, in the sense of a realist interpretation of theory, is proffered as an explanation of novel success, and the fact that realism explains novel success, while carrying further, defeasible consequences, functions as a premise in an argument for MER" (103). He goes on to note that a realist interpretation of a theory "takes the explanatory mechanisms of the theory to be representative of actual processes in nature that produce the observationally accessible effects whereby the theory is tested" (103). This is admirably clear: epistemic realism has nothing to with the explanation of success. Rather, realism, construed in ontological terms, explains success, and this allows us to infer MER. [21]

It's worth pointing out that Leplin's epistemic realism is not Laudan's epistemic realism. Recall from the previous section that Laudan's pessimistic induction is launched against what he takes to be an epistemic version of realism. I suggested in that discussion that Laudan mistakes his target; he argues against the view that the success of science requires the successful reference of the terms of science and so he argues against a semantic conception of realism. Nothing hangs on this, really; I started out by noting that one realist defends epistemic realism and one anti-realist repudiates it, and so it's worth

---

[20] See chapter five, especially pp. 102-104 and 116-135.

[21] Exactly what the benefit is of this further inference is not altogether clear, but this is something we need not explore here. Suffice it to say that epistemic realism is not the sort of realism operative in the success argument, even in Leplin's extended defense of epistemic realism.

noting that Laudan's repudiation misses Leplin's mark. No worries, though, because there is independent reason to reject the epistemic conception of realism.

1.3 the ontological conception of realism

The discussion of the previous sections has laid bare my sympathies with the ontological conception of realism. I have argued that the semantic conception is parasitic upon it, and I have argued that it is superior to both the semantic and the epistemic conceptions of realism when it comes to making sense of the success argument for realism. I want in this section to defend the ontological conception of realism and to show how it makes the success argument work.

Science is successful; we can take this to mean (at least) that its predictions pan out and that it allows us to control our environment.[22] How is it able to do this? The realist answers: 'because the entities and processes that science says exist really do exist. If this weren't the case then the theories wouldn't yield successful predictions.' A certain scientific theory states that a magnet, when broken in half, will produce two halves, each of which will attract or repel metals. We take a magnet, we break it in half, and the halves each behave in the way specified by the theory. How do we explain this? By noting that the world is as the theory says it is. Of course, the theory could be wrong- but, then, the fact that the theory made this prediction would be nothing more than a coincidence. Switching our focus from the magnets to the theory, how would we account for the fact that a false theory told us how things in the world would behave? Nothing comes to mind; it would be a miracle if a theory were able to do this. And so, the realist concludes, we can account for the success of science only by assuming that our scientific theories accurately tell us what the world is like.

On an intuitive level, then, the success argument is a straightforward argument for the literal treatment of the ontological commitments of scientific theories. But lots of smart people reject this literal construal, adopting instead the semantic or epistemic construals considered earlier. Why? In our discussion of the semantic conception of realism, we considered one such reason: the belief that talk of ontology is at bottom metaphysical and so wrong-headed. But the only reason to make a semantic ascent is to preserve the intuition that realism is a thesis about what there is. While we might need to revert to semantic *formulations* of realism, we ought not to identify realism with any particular theory of reference or truth. As I suggested earlier, the claim that the world behaves in such-and-such a manner will not become false if the proper analysis of truth (or reference) turns out to be a non-correspondence theory (or a non-causal theory).

A different objection might go as follows: "what explains the success of a theory is the fact that the theory accurately represents the world, not merely the existence of the theory's postulated entities. While the thesis of realism can be stated in a way that is, on the surface, purely ontological, it must be construed in semantic terms in order to make sense of its role in the explanation of scientific success." The example given above is rigged, the objector might say: you can cast it in ontological terms, but what really must be going on is that the theory must include statements like 'if a magnet is broken in half then each of the halves will itself be a magnet' and 'a magnet attracts and repels metals',

[22] I'll have more to say about this in the next chapter.

and these sentences must contain terms whose reference is specified in such a way that they actually refer, and the sentences must have truth conditions that are realized just in case the terms employed in them refer in the right ways.... Without all of this, you don't have an example of a theory yielding a successful prediction; you've just got some stuff happening in the world.

This objection is apt because it is keyed to the success argument. But is it right? What explains the success of a theory[23] is the fact that the world contains the entities postulated by the theory and works according to the laws stated by the theory. Sure, this talk of postulation and the like is semantic. What explains the theory's success is the fact that the representation of the world offered by the theory is accurate. But does this make realism semantic? Only in the trivial sense that everything is semantic. Sure, the terms of the theory of magnetism must refer if the theory is to get off the ground. But the terms of every theory must refer in this way. What makes the theory of magnetism a good one -- what makes it a *successful* one -- is that the things the terms of the theory purport to refer to really do exist, and they really do behave in the manner specified by the theory.

Any statement about the world -- even an unsuccessful theory -- must contain terms that purport to refer to the world. But such reference is insufficient to account for success. My sentence 'if you rub two oak leaves together, you will produce a spark' contains genuinely referring terms, but it is not a successful theory. What would be required for it to be successful would be oak leaves behaving in just such a way- it would be successful just in case it really were the case that rubbing oak leaves together produced a spark. Of course, this isn't the case, and that's why the theory is a bad theory: it is an account of the behavior of the world (it's terms genuinely refer to things in the world) that gets their behavior wrong. Genuine reference allows our theories to be about the world, but it doesn't make them successful. Something else -- something extra-semantic -- is required for that.

Perhaps this extra-semantic something is reference to mechanisms. We can expand our ontology to include mechanisms with entities, and we can say that a successful theory is one that genuinely refers to mechanisms. Since there is no such mechanism operating on oak leaves, the theory is not a genuinely referring one, and we can explain it's failure on these grounds. Well, we could say this. But why on earth would we want to? The motivation for adopting a semantic version of realism is to avoid doing ontology, and here we would be arguing for a complication of our ontology in order to save our semantic alternative to it. Besides, this move would only prolong the inevitable. For let's suppose that we do expand our ontology to include mechanisms. Then the semanticist is going to point to the non-reference of sentences postulating mechanism M in order to explain the failures of theory Tm. And the ontologist is going to point to the non-existence of M in order to explain the failure of theory Tm. The debate will be an exact replica of the one rehearsed earlier: the ontologist will argue that reference fails because the thing doesn't exist, that non-existence underwrites non-reference, and that talk of

---

[23] For the realist, anyway. Recall that the aim of this chapter is to develop the strongest statement of realism possible, so I am giving the realist all he needs. Whether the explanation of success succeeds will be considered in chapter 4.

reference only makes sense against this ontological backdrop. The appeal to mechanisms takes the same argument and just pushes it back one step.

We can grant, then, that there is a semantic tinge to scientific realism -- our theories involve theory-world relations and so must be about the world -- but this tinge doesn't explain success. (To put the point another way, realism presupposes semantic realism, but semantic realism is insufficient to account to success.) We need to go beyond representation and find the entities and mechanisms in the world that behave in the manner specified by our theories, and this will explain why the world has behaved in the way the theories said it would. The most semantics can do is refine our discourse about this talk of discovery. If you've got qualms about doing ontology, then revert to your refinements and talk about the truth conditions for statements about the mechanisms. Those of us with stronger constitutions will continue to ask whether the world is really as our successful theories say it is, and we will still wonder how else that success might be explained, if it is not.

1.4 conceptions of anti-realism

If the arguments of the preceding sections are correct, then realism is the thesis that (some or most of) the entities and processes postulated by science exist. This is by no means a complete account of realism: one can press it on the force of the qualification 'some or most', and one can argue that the ontological commitment of the realist be limited to just entities ('entity realism') or that it should include processes ('theory realism').[24] A full-fledged defense of the ontological account of realism would have to take account of all these objections. Since I am interested not in defending such an account, but in defending the role of the ontological account in the success argument, I needn't go this extra mile. Suffice it to say that the arguments that have been offered above show that the success argument is properly understood in terms of the ontological conception of realism.

But we are not only interested in figuring out how the realist should formulate the success argument; we are also interested in whether the anti-realist can account for success. So let's have a look at the types of scientific anti-realism available to us. The most common way of slicing up the terrain is as follows. (In this, I follow Michael Devitt: 1997, 14-22). Realism has two dimensions: an existence dimension (according to which the things that science says exist really do exist) and an independence dimension (according to which the things that science says exist, exist independently of the mental lives of humans), and the anti-realist denies one or the other of these dimensions. Let's call any anti-realism that denies the existence dimension a *strong anti-realism*, and any anti-realism that denies the independence dimension a *weak anti-realism*. So the idealist and the constructivist are weak anti-realists, while the instrumentalist and the constructive empiricist are strong anti-realists[25]. The weak anti-realist must argue that the success of science can be explained without appealing to the mind-independent existence of the

[24] The distinction, of course, is Hacking's. It's worth noting Goodman' commonplace that entity realism is "not worth fighting for"- it is too minimal (in Devitt's apt phrase, it is anti-realism with a fig leaf"). The interesting sort of realism is theory realism, which says not just that the things exist that science says exist, but that the world behaves according to the laws that science states. It's worth noting also that some realists (notably, Cartwright and Hacking) think theory realism is false.

entities and processes postulated by science, and the strong anti-realist must argue that the success of science can be explained without appealing at all to the existence of these entities and processes.

Historically, neither sort of anti-realism has been able to get off the ground. The weak anti-realisms of idealism and constructivism have been unable to construct an adequate picture of science (in the case of constructivism) or of metaphysics (in the case of idealism). Kuhn's constructivism, for instance, requires that we treat different periods in science as speaking to wholly (and literally) different worlds- a result that flies in the face of our abilities to understand and compare those theories. Idealism doesn't fare much better: it requires us to accept the unreasonable conclusion that unperceived things do not exist, thereby robbing us of our rich history of discoveries. The failures of the strong anti-realist program of instrumentalism are equally well-known: the failure of the verificationist criterion of meaningfulness served to virtually annihilate the positivist project, and instrumentalism went with it. Constructive empiricism, as it is presented in the writings of Bas van Fraassen, seeks to raise that ship by serving up a slightly different version of instrumentalism. Whether it succeeds in this is, of course, another matter altogether. (We will look at it in some detail in chapter three, where I argue that existing scientific anti-realist programs fail to explain the success of science.) Suffice it to say for now that neither weak anti-realism nor strong anti-realism is widely regarded as a viable philosophy of science. But these are regarded as the only possible anti-realist philosophies of science, and so it is thought that there can be no anti-realist philosophy of

[25] Well, the constructive empiricist wants merely to remain agnostic about the existence of unobservable entities, and so doesn't deny their existence, *per se*. Qualification noted; the view is stronger than it is

science that accounts for the success of science. This is what gives the success argument for realism it's bite.

There are anti-realisms other than the strong and weak versions I have identified (Dummett's, for example). But these others types of anti-realism presuppose a non-ontological conception of realism, and so they are beyond the scope of this analysis. The task before us has been to map that portion of the philosophic landscape relevant to the success argument for realism. We have done that, and we have identified a single sort of realism and a couple of sorts of anti-realism that respond to realism, so understood. We must now turn our attention to evaluation of the argument for this sort of realism (and the concomitant arguments against anti-realism); but first, let us look a bit more critically at the concept of success.

---

weaker.

2. SUCCESS

The realist says that only scientific realism can account for the success of science. In the previous chapter, I assumed that the relevant sort of success is explanatory and predictive accuracy and environmental control. I didn't do much to support this assumption, other than a brief parenthetical remark to the effect that this is how Putnam (1975) conceives of success. Given my stated goal of analyzing the success argument for realism, it would seem rather important that we give the notion of success its fair due. We know what realism is; we should now endeavor to know what success is.

Putnam (1975) is the canonical statement of the argument from success. He frames the argument as follows: "That science succeeds in making many true predictions, devising better ways of controlling nature, and so forth, is an undoubted empirical fact. If realism is an explanation of this fact, realism must itself be an overarching scientific hypothesis" (141). So there's prediction and control of nature; we can add to this list explanatory and methodological success. The question we must ask is, 'which conception of success stands most in need of explanation, and which conception lends itself most readily to explanation in realist terms?' This will allow us to construct the realist's success argument in the strongest possible terms. What I propose to do in this chapter is examine each type of success and consider its role in the success argument. As before, my intention is not to promote anti-realism or realism. Rather, I seek to determine whether there is a cogent enough conception of success to underwrite the realist's argument.

2.1 instrumental success

Putnam never really follows up on the conception of success as manipulation or control of nature, nor do any of the other proponents of the success argument.[26] In conversation about relativism, one frequently hears things like: 'if mysticism [e.g.] is really as good as – just different than – Western science, then why can't the mystic build a rocket or develop a vaccine?' The assumption is that mysticism cannot rival the technological advances of Western science, and this is reason to think Western science provides us with certain truths about the world that mysticism fails to grasp. Slightly beneath the surface is a principle like 'if a system of thought allows you to manipulate your environment, then this is reason to think that system accurately reflects the structure of that environment.' Abstracting from the context of relativism, we get a success argument for realism: 'Western science has produced vaccines, vaccines work, vaccines are based on the germ theory of disease, this is evidence that the germ theory of disease is true.'

The rhetorical force of this version of the success argument is nearly overwhelming. But it is doubtful that it constitutes a substantive conception of success. For what is it to achieve instrumental success, except to usefully employ explanatory and predictive success? To develop a vaccine is to just to put to work the several explanations and predictions afforded by the germ theory of disease. The force of this manner of argument is considerable when cast in the context of relativism, for the realist (read: defender of Western science) can ask for the comparable applications of non-Western explanations and predictions. But if we limit ourselves to arguments for realism, not arguments against relativism, appeals to technology do little original work.

In any event, this is how I propose to dispense with the conception of success as instrumental control: it reduces to predictive and explanatory success. Let us, then, proceed to those conceptions of success.

2.2 explanatory success

Another sort of success that a theory might enjoy is explanatory success. The idea here is simple: some observation, event, or result has been heretofore unexplained (or unsatisfactorily explained), a theory comes along and accounts for the phenomena, and the theory is therefore regarded as true. Explanatory success isn't as impressive as predictive success – the mere fact of accounting for phenomena known to exist doesn't have as much probative force as accounting for something heretofore not known to exist – but it is still intuitively plausible to claim that, if a theory offers the only plausible account of some phenomena, then that theory is likely true.

This intuitive formulation needs clarification. There is, to begin, the vague notion of a theory offering 'the only plausible account.' The mere offering of an account is, of course, insufficient to establish the truth of a theory: 'there was a lunar eclipse because the sun ate the moon' offers an explanation, but this doesn't lead us to accept its ontology. When we say we want a plausible explanation, then, what we want is a good explanation.[27] Moreover, we want this explanation to be the only game in town: if there

[26] Ma. Levin (2000) is a notable exception.
[27] 'Good' is intentionally vague. Most people would argue that explanations need to be true – that false theories cannot explain – and so would contend that 'good' ought to be replaced by 'true'. The problem is immediately clear: if we replace 'plausible' with 'good' and 'good' with 'true', then the argument for

are two possible explanations of some event, and if the explanations have equally acceptable but incompatible ontologies, we have to decide between the explanations before we can infer the truth of one of them.

Given these constraints, one question we can ask is whether the inference from explanation to truth is warranted. Another question is, 'can this sort of success underwrite the realist's claim that, absent realism, there is no accounting for the theory's success'?

Peter Lipton suggests that explanations can be trumped up and so are not reliable indicators of truth: "When data need to be accommodated, there is a motive to force a theory and auxiliaries to make the accommodation. The scientist knows the answer she must get, and she does whatever it takes to get it. " (Lipton 1991: 140). A scientist who seeks to accommodate existing phenomena, Lipton contends, already knows what results she is seeking to achieve, and this influences her design of both theory and experiment. Hence explanations can speak to particular situations, but are not therefore more likely to be true than other theories. (Here we can adduce the cases from the historical induction as examples of explanatory, false theories). Not only is there little reason to think that the truth is the *only* way to account for a theory's explanatory success; there is little reason to even regard explanatorily successful theories as true.[28]

---

realism is question-begging: 'if a theory offers the only true account of some phenomena, then that theory is likely true.' There are ways around the problem, to be sure, but I think it is best at this point to not complicate matters and stick with my vague formulation. We will have opportunity to examine the issue in a few pages and, again, later on in the essay.

This leads to a related problem. Rarely is it the case that an explanation can be neatly tested, and what constitutes a good explanation is not often easily determined. Often, a theory will be regarded as a good one, only to be rejected later- that is, the explanatory success of a theory will be treated as reason to attribute truth to the theory, but the theory will later be proven false. Hence there is no inferential path from explanation to truth; explanatory success cannot underwrite the success argument for realism.

The realist will likely demur: false theories cannot explain, so there is no threat to realism. While some theories that are thought true might be thought to explain, they cannot actually explain unless they are in fact true. Hence the apparent counter-example fails: the fact that false theories are sometimes thought to be explanatory is no reason to deny that the (actual) explanatoriness of a theory is reason to think it is true.

As suggested above, this argument seems to make trouble for the success argument for realism. If truth were guaranteed by explanation, then the realist would stand accused of begging the question: he would be arguing that, if a theory explains then the theory is true, and only true theories can explain. This will convince no one; truth has got to come from the theory, not the theory of explanation.[29]

---

[28] The account of success I sketch in chapter six of this essay will sound a version of this theme.
[29] Michael Levin suggests that the realist can avoid the problem by simply "redescrib[ing] the datum. Instead of saying that what is to be explained is the explanatory success of science – which does assume realism – he can say that what needs explaining is the *apparent* explanatory success of science. We have theories that predict phenomena, would unify old phenomena if they were true, etc. How else account for these wondorous properties unless those theories *were* true. [sic]" The suggestion is that we re-identify the *explanans-* what the realist seeks to account for is not the explanatory success of particular theories (if explanations must be true, this *would* beg the question) but, rather, the fact that scientific theories appear to be explanatorily successful. Hence we iterate the success argument: given that theories must be true in order to explain, and given that scientific theories often appear to explain, scientific theories must be true, or else their apparent success would be a miracle.

But didn't we spend a good deal of time in the last chapter dismissing the view that realism has to do with truth? And doesn't the argument of the previous paragraph presume that realism consists in a commitment to the truth of scientific theories? Well, yes- but this is no reason to suspect either my rejection of the semantic construal of realism or my response to the objection considered above. I have been considering the claim that only true theories can explain, and I have been arguing that this view of explanation cannot legitimately by adduced by the realist in support of realism. The reason for this is that the success argument for realism must regard the truth of scientific theories (that is, the accurate representation by those theories of reality) as the conclusion of an argument that takes as its central premise the role of realism (the view that scientific theories accurately represent reality) in explaining scientific success. And I have been belaboring the point that it would be a cheat for the realist to sneak truth (accurate representation) into the picture by packing it into an account of explanation. If you regard realism in semantic terms, then this conclusion should be clear. If you regard it in ontological terms then the conclusion is a bit more distant, but still pretty clear: for scientific realism to have any bite, it must allow science to do the work, not the theory of explanation.[30]

This move will work, but notice (as Levin points out) that it reduces the argument from explanatory success to the argument from predictive success. How are we to account for the fact that, generally speaking, scientific theories appear to explain? By assuming that realism is true. But in doing this we are accounting for the fact that our theories will be apparently (explanitorily) successful, rather than explaining their actual explanatory successes. Hence Levin's re-formulation reduces the argument from explanatory success to the argument from predictive success.

[30] This is reminiscent of Fine's view (to be considered later): science, not philosophy explains scientific success. Perhaps this is his point: neither philosophy nor the theory of truth can win the day for realism- science must do that job.

We have been considering objections to the view that a theory's explanatory success is a reason to be a realist about that theory. Ernan McMullin (1984: 26-7) advances an argument to the contrary: "scientists construct theories which explain the observed features of the physical world by postulating models of the hidden structure of the entities being studied. This structure is taken to account causally for the observable phenomena, and the theoretical model provides an approximation of the phenomena from which the explanatory power of the model derives." He proceeds to give examples (from geology and chemistry) of theory change and growth that have been continuous and fruitful. These successes, he contends, stand in need of explanation, and only realism will do. While Leplin would have us explain how it is that a theory can predict some novel result, then, McMullin would have us explain how it is that a structural explanatory hypothesis can yield an increasingly accurate account of the physical world.

Notice that the sort of explanation McMullin is concerned with is not explanation of events. Such explanations are the kind that seem too easily trumped up, for there is a discrete phenomenon to be accounted for, and the theory is designed to do that job. Structural explanations are not designed to account for some particular event; rather, they postulate unobservable structures that are responsible for the behavior of the observable world, and what is in need of explanation is the fact that "the growth in our knowledge of structure has been relatively steady" (27). What McMullin is after, though, is more an account of predictive success than an account of explanatory success: what stands in need of explanation is the remarkable rate at which our predictions – based on structural postulates – of the behavior of the observable world come out true. Hence, while structural explanation is, unlike the sorts of explanation considered earlier, sufficiently robust to underwrite the success argument, it underwrites the predictive success argument, not the explanatory success argument.

In the final analysis, then, there is no robust argument from explanatory success for realism. The only non-question begging version of the argument reduces to the argument from predictive success (as does McMullin's version of the argument). Additionally, the malleability of explanation makes it unlikely that there is a connection between it and truth: theories and data are retrofitted, as it were, to fit with one another. While the argument from explanatory success remains intuitively plausible, it seems unable to sustain critical analysis.

2.3 predictive success

A popular thing to think about success is that accurate prediction is the best indicator of the truth of a theory.[31] The matter is usually cast in terms of a debate over whether prediction or explanation is the better indicator. The argument for the preference for prediction goes like this: while we might be able to explain away a theory's accommodation of known facts by appeal to *ad hoc* manipulation of data or experimental situations, it is far less likely that a theory could be rigged to predict future events. If no one knows that there's going to be an eclipse, and the theory says that there will be one, and there is one- well, this is reason to think the theorist knows something we don't.

[31] I'm using "truth" with a small 't' here: I'm not advocating a semantic conception of truth as much as I am using 'true theory' as shorthand for 'theory that accurately represents the world.'

One problem with this conception of success is that 'true prediction' is a rather vague term, and we may fairly ask just what is needed for a true prediction to constitute success. Consider the prediction: 'there will be an eclipse.' This is fairly specific, even though there is no temporal constraint involved. What if the content of the prediction were less specific, though- something like 'there will be planetary movement', or 'something will happen involving the sun and the moon' or 'something will happen.' We would be inclined, I think, to regard the first and third of these as too vague to count as acceptable, while we might be inclined (maybe) to accept the second. The content of the prediction matters. So does the time-frame: if the prediction is 'something will happen between 10:40 and 11:00 on the 19th of October', we would be likely to accept it, while 'there will be planetary movement in the next 500 years' remains intuitively unacceptable. The problem, in short, is that we don't know what counts as a 'true prediction'. There seem to be needed constraints on both subject matter and modulus, but no one seems to know what those constraints are.

Another difficulty for prediction is introduced by Peter Lipton (1991: 133-157), who argues for the comparative advantage of prediction to explanation ('accommodation') in stating the success argument for realism. Lipton argues that prediction lends more inductive support to theories than does accommodation of known facts- that the predictive success of a theory is more reason to believe the theory than explanatory success.[32] In so doing, he raises an objection that poses a problem for the conception of success as predictive success. To illustrate the problem, Lipton asks us to imagine twin

[32] This sounds like an endorsement of the epistemic conception of realism. As before, though, (see previous note) it is only a convenience of exposition.

scientists who, independently of one another, construct the same theory. The only difference between them is that one of them accommodates a given result while the other predicts it. If prediction provides more reason to believe a theory than accommodation, Lipton notes, then the predictor would have more reason to believe the theory than the accomodator. But the theories are exactly the same, and one would think that there is therefore equal reason to believe the one as to believe the other.[33] What Lipton's twins case seems to show is this: if predictive success provides reason to believe that a theory is true, and predictive success is an indicator of realism, then all one needs to do in order to establish realism is know nothing at all.

Lipton proceeds from this argument to consider and reject some defenses of (the probative force of) prediction and then advance his own defense of the preference for prediction. But he never dismisses this argument against treating prediction as a reason to believe, *viz*., that on this view, theoretical support is attained through ignorance of the implications of one's theories. Perhaps prediction is preferable to accommodation in analyzing success; if so, so much the worse for accommodation: the problem with prediction remains. Indeed, it has rather serious normative implications (scientists that want to be successful should learn as little as possible about the theories they employ) and it dulls the distinction between predictive success and explanatory success- the only difference between them is the psychological state of the theorist, and this is hardly sufficient to underwrite a difference of probative force. Indeed, if we conceive of realism in ontological terms, there is absolutely no relative difference between predictive success

[33] Collins (1994: 215-219) points to a similar problem: granting prediction an epistemic priority over explanation yields counter-intuitive results.

and explanatory success. In short, given Lipton's analysis, the coincidence of prediction and result (the coincidence that underwrites the (predictive) success argument) is, in fact, more manipulable than imagined: simply stay ignorant of the implications of your theories until you test them, and then claim predictive success.

Underlying Lipton's analysis is the assumption that the key difference between prediction and explanation is the psychological state of the experimenter: the scientist who explains has some bit of knowledge that is lacked by the scientist who predicts. We might avoid Lipton's problem by rejecting this psychologistic reading of prediction. This is what Jarrett Leplin (1997) does: he contends that a successful theory is one that "uniquely explain[s] and predict[s] an observational result without itself depending on that result for its content or development" (64). It's not merely prediction that provides reason to believe, he thinks- it's prediction of novel facts. And what makes a fact novel for a theory is not that the theorist does not know about the fact, but that the result is independent of the theory and unique to the theory. For any theory T and any observation result O, O is independent of T just in case "there is a minimally adequate reconstruction of the reasoning leading to T that does not cite any qualitative generalization of O" (77). And O is unique to T just in case "There is some qualitative generalization of O that T explains and predicts, and of which, at the time that T first does so, no alternative theory provides a viable reason to expect such instances" (77).[34]

[34] "By *the qualitative generalization of O*," Leplin writes, "I mean *the effect itself*, the type of phenomenon -- itself a type -- that O instantiates, independent of considerations of quantitative accuracy" (73).

The independence condition ensures that temporal constraints need not apply; past events can be novel for later theories (like the precession of the perihelion of Mercury was for Einstein's theory of general relativity). And the uniqueness result ensures that there is reason to attribute truth to the theory: if other theories also predict the event, then there is no more reason to attribute truth to one theory rather than the other. Taken together, these conditions operate on predictions to avoid the problems considered above: it isn't enough to be merely ignorant of the implications of one's theory to make those implications novel; rather, consideration of the implications must not essentially enter into the development of the theory. On Leplin's analysis of Lipton's twins case, then, either both results would be novel or both would not be novel, regardless of differences in the theorists' knowledge.

On the face of things, then, it looks like Leplin's account of novel prediction serves well the realist cause. If a theory predicts some fact that in fact occurs, and if we can't attribute this success to another theory, and if the result doesn't figure into a reconstruction of the reasoning that led to the theory, then this is reason to think that the theory has latched onto some genuine feature on the world.[35] We'll have opportunity later to critically examine the argument from novel predictive success to realism; suffice it for now to say that Leplin's analysis is a likely candidate for the job of filling out the success argument.

---

[35] There remain the problems, noted at the start of this chapter, concerning the modulus and specificity of predictions. These problems remain but do not seem sufficiently serious to warrant a rejection of predictive success as probative. I would submit that the rough analysis given in that discussion is sufficient to settle

2.4 methodological success

Richard Boyd thinks the sort of success that needs explaining is methodological success. In a series of papers (1980, 1984, 1989, 1990, 1992), he argues that realism is shown true by its unique ability to account for the "instrumental reliability of scientific methodology" (1984: 58). He argues that the scientific method consists in a dialectical method that yields convergence of theories on truth, and he sets out to determine what it is that makes this possible. His answer, of course, is scientific realism.

It's worth noticing how very different Boyd's conception of success is than the others. Predictive success and explanatory success are properly attributed to theories. Methodological success is a property of the scientific method. Boyd is impressed with the predictive and explanatory capabilities of scientific theories, but what really impresses him is our ability to keep coming up with theories that exhibit these properties. Boyd (1980: 626) explains this in more detail:

> By the 'instrumental reliability' of a scientific theory I will mean its ability to provide (given suitable 'auxiliary hypotheses') approximately accurate predictions about the behavior of observable phenomena. By 'instrumental knowledge' I will mean the knowledge about particular theories that they are instrumentally reliable, and the concomitant knowledge about observable phenomena. By the 'instrumental reliability' of methodological principles, I mean their capacity to contribute to the production of instrumental knowledge.

He goes on to wonder how to account for such instrumental reliability, and he argues (predictably) that only realism will do.[36] So what needs explaining, for Boyd, is the

---

the matter- we have a sense of what standards a prediction needs to meet to be probative, even if we do not have a precise statement of those standards.

[36] Actually, Boyd is concerned in his (1980) to defend a naturalized epistemology of science, according to which the background beliefs against which new theories are tested are assumed to be true, and this

ability of scientific methodology to produce the knowledge about theories that they are predictiively fruitful. Alternatively, he wants to explain how it is that the methodological principles of science allow us to choose theories that will yield accurate predictions. Science is methodologically successful, then, in virtue of the fact that it leads us to adopt predictively successful theories.

A little reflection reveals that explaining this sort of success is a serious problem- one that is related to, but distinct from, the problem of accounting for theories' abilities to yield accurate predictions. Boyd (1980: 618) poses the problem rather starkly:

> Given any plausible initial total science there will be infinitely many equivalence classes under the relation [of observational equivalence]. One way of putting the problem of the instrumental reliability of scientific method is this: In fact we choose one of these equivalence classes each time we accept a theory, and we do so on the basis of finitely many observations. So some other criteria other than consistency with observational data *must* be at work. Call these the *'extra-experimental'* criteria. Whatever these extra-experimental criteria are, they work. In the long (but not very long) run we get quite good predictive theories. Why do these criteria work?

Arthur Fine (1984: 88) is also keen to this problem. He calls it the problem of the 'small handful':

> At any given time, in a given scientific area, only a small handful of alternative theories (or hypotheses) are in the field. ... Moreover, in general, this handful displays a sort of family resemblance in that none of these live options will be too far from the previously accepted theories in the field .... Why? Why does this narrowing down of our choices to such a small handful of cousins of our previously accepted theories work to produce such good successor theories?

---

assumption grounds the presumption that new theories -- which extend the old ones -- are similarly true. This is not unlike the irrealist account of success I consider in chapter six.

Boyd, unlike Fine, casts the problem in the shadow of under-determination; this is because Boyd regards under-determination as the most formidable threat to realism. But there is nothing wedding the problem to that threat; we can abstract from it to ask what it is (in addition to the finite number of observations we make) that underwrites our theory choice, and we can continue to marvel at the fact that, whatever it is, this extra-experimental criterion allow us to pick such good theories. How is it that, given the panoply of theories available to us, and given the fact that we consider only a handful of these theories, we manage somehow to settle on the ones that work?

I trust that these remarks suffice to differentiate methodological success from predictive success; no account of the latter is going to serve as an account of the former as well. Yet it is undeniable that science is methodologically successful. Thus we have two distinct conceptions of success. We must ask ourselves which can be of greater service to the realist. Must realism be true – must the entities and processes postulated by scientific theories exist – in order for science to be methodologically successful? Not obviously. The connection between methodological success and realism is parasitic on the connection between predictive success[37] and realism.

What we want is an account of the instrumental reliability of scientific method- of the capacity of that method to produce successful theories. Suppose realism weren't true- how could we account for the fact that we consistently produce theories that are predictively successful? Boyd's answer is that we couldn't; hence his (methodological) success argument for realism. Thus stated, though, the connection between instrumental reliability and realism hinges on the connection between prediction and realism. There are two steps to Boyd's argument: (i) our theories are predictively successful, and (ii) we consistently produce such theories. Boyd is impressed by our ability to consistently produce theories that work. He concludes that the scientific method is a reliable method for producing theories that tell us what the world is like.[38] Alternatively, he concludes that we couldn't consistently produce predictively successful theories unless our method for producing theories was a method for producing theories that tell us what the world is like. What's doing the work is the predictive success of the theories- our methods for producing theories must be successful because our theories are successful, and there is no way to account for this except realism.

Boyd is attempting to make sense of the claim that the predictive success of scientific theories is reason to think those theories are true (that is, he is, *pace* Laudan and Fine, trying to account for the abductive connection between success and realism), and the reason he offers is that the general success of science gives us reason to think that the scientific method gets at the truth. Put another way, Boyd is accounting for the abductive connection between predictive success and realism by adducing a second abductive connection- this one between the general predictive success of science and realism. That this seems, on its face, to be question-begging needn't concern us here[39]; the point is that

---

[37] Given the analysis of preceding sections. I will hereafter talk only about predictive success.

[38] This epithet – 'tell us what the world is like' – is rough, but it allows for a relatively fluid exposition of the argument. I don't intend to saddle the realist with undue commitments: I simply leave out all the usual qualifiers ('more or less', 'some or most of the entities', etc.) for ease of exposition.

[39] The short answer is that Boyd argues for a naturalized realism- realism is an empirical hypothesis that seeks to explain the success of science, and the meta-inference to methodological success is "part of an overall 'realist package'. This package should be compared as a whole with the overall alternatives. If the realist package is philosophically more defensible overall, then the charge of circularity falls away" (Pappineau 1996: 15-16).

Boyd's methodological success gets its bite from the predictive success of science: he takes as fundamental the claim that science is predictively successful and that realism accounts for this, and he accounts for the fact that realism accounts for predictive success by appealing to the methodological success of science.

To be sure, methodological success does not reduce to predictive success: that our theories are predictively successful is one thing; that we consistently select predicticively successful theories is another. But – and this is the point – there is no inferential path from methodological success to realism that does not trace the path from predictive success to realism. There can be no methodological success argument for realism that does not presuppose the predictive success argument for realism. Hence the methodological conception of success does not underwrite the success argument for realism.

3. ANTI-REALISM

Recall that Putnam's (1975) invocation of the success argument for realism was fueled by the contention that there are no viable anti-realist philosophies of science to be had. This combined with the need for an account of the success of science to make realism the default position in the philosophy of science. In the intervening years, however, anti-realists have been hard at work. Bas van Fraassen has developed a sophisticated neo-positivistic philosophy of science that he calls *constructive empiricism*. Arthur Fine has developed a naturalized philosophy of science that he calls the *Natural Ontological Attitude* (*NOA*). Constructive empiricism is by all accounts anti-realist. NOA is, by Fine's account, neither realist nor anti-realist. The present chapter offers an analysis of these alternatives to realism.

I argue that neither NOA nor constructive empiricism offers a serious alternative to the realist's explanation of scientific success. While Fine and van Fraassen offer ingenious (and sometimes incisive) criticisms of realism, neither of their positive programs is able to endure sustained critical analysis. Fine's Natural Ontological Attitude is shown to be a hybrid version of realism about ontology and anti-cumulativism—a sort of weak epistemic realism that is unable to account for the success of science. van Fraassen's position is a genuine anti-realist alternative to realism, but it hasn't the resources to generate an informative explanation of success.

In the final analysis, the anti-realists seems to be without a response to the realist. While she can offer criticisms of realism, the anti-realist cannot proffer an explanation of scientific success to compete with the realist's. Even if the realist's account is flawed, he's got one, and this is enough to tip the scales of intuition in the realist's favor.

3.1 Fine's *Natural Ontological Attitude*

Arthur Fine thinks that both realism and instrumentalism commit the sins of inflationism. Indeed, he thinks that any party to the realism debate must commit this sin, simply by virtue of the adoption of a philosophical program. The only way to avoid inflationism is to avoid philosophical accounts of science. Fine claims to do this with his NOA. He argues that the realist and the anti-realist share a common ground, called the core position, and that they differ only in the philosophical baggage they attach to it. The trouble with the realism debate, Fine thinks, is that both the realist's and the anti-realist's baggage are metaphysically and epistemologically excessive. NOA endorses the core position, *simpliciter*, and thus is supposed to capture all that is true in both realism and anti-realism without committing the mistakes made by either of those positions.

Fine describes this core position as one that "accept[s] the certified results of science as on a par with more homely and familiarly supported claims" (1984: 96). To accept the results of science in this way, Fine tells us, "is to take them into one's life as true, with all that implies about adjusting one's behavior, practical and theoretical, to accommodate these truths" (1984: 95-96). The realist and the anti-realist both accept the core position, Fine says. But the anti-realist adds to it "a particular analysis of the concept of truth, as in the pragmatic and instrumentalist and conventionalist conceptions of truth. Or the anti-realist may add on a special analysis of concepts, as in idealism, constructivism, phenomenalism, and in some varieties of empiricism." And the realist adds to the core position "a desk-thumping, foot-stamping shout of 'Really!'" (1984: 97). The realist endorses the core position but adds to it a metaphysics- 'everything science tells us is true, and the world really and truly is like that.' And the anti-realist endorses the core position but adds to it a metaphysics – 'everything science tells us is true, but the world isn't really like that...' – and an *a priori* epistemology – 'because 'the reach of our knowledge about the world extends only to the observable.' Fine thinks that neither of these philosophical add-on's contributes anything of substance to the core position; they are attempts to force science into the pre-existing philosophical molds.[40] In developing NOA, he seeks out a position with nothing added.

The NOAer accepts the claims of science without attaching to them any philosophical significance. They are not held to have any ontological import (*pace* the realist) or any special non-literal interpretation (*pace* the anti-realist). Fine seeks to clarify this by contrasting NOA with realism and instrumentalism:

> The point is that realism requires two distinct elements. It requires belief and it also requires a particular interpretation of that belief. Thus anti-realism, in particular instrumentalism, pursues the following strategy. It does not withhold belief, then it offers instead a non-realist interpretation of that belief .... But the reader will no doubt notice that there is an interesting third way. For one can go along with belief but then simply not add on any special interpretation of it- neither realist nor anti-realist. That is the way of NOA (1986: 176).

NOA, then, resists the move from belief in a theory to attribution of truth to the theory: "NOA is inclined to reject all interpretations, theories, construals, pictures, etc., of truth,

[40] "These extra-scientific orientations [of realism and anti-realism] to science preserve some cherished elements of recognizable philosophical schools. In each case, inflationism is the consequence of trying to

just as it rejects the special correspondence theory of realism and the acceptance pictures of the truthmongering antirealisms" (1996a: 149).[41] In the place of these theories of truth, "NOA holds a 'no-theory' conception of truth" that "accepts the usual logic and grammar of truth, including its redundancy property." Lest the reader be confused, Fine makes clear that "This no-theory attitude towards truth separates NOA from realism, since realism is committed to a special interpretative stance" (1986: 175).

So everyone is expected to believe that the sentences of (the mature theories of) science are true. The question is, 'what does this belief amount to?' (In Fine's terms, 'how are we to interpret this belief?') For the realist, it amounts to the endorsement of a metaphysical picture of reality: we interpret the belief literally. For the anti-realist, it amounts to the endorsement of the scientific picture, but not as a blueprint of the universe: we interpret it non-literally. But what does acceptance come to when it comes without a philosophical interpretation? Whatever the merits of the realist's and anti-realist's theories of truth, those theories commit them to well-defined positions. But what sort of position does the no-theory theory of truth commit you to? Fine thinks that it commits him to something unlike either the realist's or the anti-realist's commitments. In order to make sense of NOA, and in order to assess NOA's ability to explain success, we need to figure out how this is supposed to work.

reconcile science with the special interpretative stance of a particular school" (Fine 1986: 171).

[41] Fine distinguishes between truthmongering anti-realisms – anti-realist programs that focus on the analysis of truth – and other kinds. I don't think the distinction plays a central role in the development of

3.1.1 NOA

I take it as obvious that NOA is distinct from those anti-realist programs that treat truth non-literally. Clearly, to accept the usual grammar and logic of truth means not treating it non-literally. And the concomitant literalism about reference ensures that NOA is distinct from idealism and other anti-realisms. So there is little danger of mistaking NOA for anti-realism.[42] But there is some confusion on NOA's other border, for the difference between NOA and realism is rather obscure. It seems, on the face of things, that to accept the usual logic and grammar of truth is to take a particular attitude towards truth- namely, a literal one. And this seems to be precisely the attitude that the realist takes towards truth. To ask the question squarely, what special interpretative stance does realism take that NOA does not also take?

One answer to this question locates the difference in the fact that the realist treats the propositions of science as correspondence-true, while the NOAer refuses to adopt a theory of truth. This is certainly the answer that Fine seems to endorse: recall his claim that the "no-theory attitude towards truth separates NOA from realism." The problem with this understanding of Fine's position – let's call it the semantic interpretation of NOA – is that it is too weak to yield the required distinction of interpretative stances. By Fine's own admission, "When NOA counsels us to accept the results of science as true, [...] we are to treat truth in the usual referential way, so that a sentence (or statement) is

---

Fine's position. Even if it does, though, the distinction won't come into play in my analysis. For interesting discussions, see Knezevich (1989) and Fine's (1989) response.

[42] It's worth noting that Ernan McMullin (1991: 106) thinks Fine is an anti-realist "simply because [he] sides with the instrumentalist critique of all forms of realism." He acknowledges, though, that "the spirit of it is different." Surely McMullin is talking loosely here, and I won't pursue the matter further than to observe that two distinct positions can be opposed to the same doctrine.

true just in case the entities referred to stand in the referred-to relation. Thus, NOA sanctions normal referential semantics..." (1984: 98).

So the NOAer does take an interpretative stance: he interprets the claims of science literally. And he must do so, for a non-literal interpretation would incline NOA towards anti-realism. What is at issue here isn't the question of whether one takes an interpretative stance, but whether the stance one takes forces one to do metaphysics. Fine seems to think that the interpretative stance required by the no-theory theory of truth is metaphysically toothless, while the stance required by the correspondence theory of truth foists upon the realist all manner of metaphysical commitments.

But this just seems wrong. The interpretation required by the no-theory theory of truth is identical to the one required by the correspondence theory of truth, and NOA accordingly collapses into realism. This can be made clear by considering a standard theoretical sentence (say, 'there are massive particles'). One can, with the realist, take the claim to be a true claim about the world. This would, given standardly accepted criteria, commit us to the existence of such particles. If we accept 'the ordinary grammar and logic of truth, including its redundancy property', as NOA counsels, we would seem to be as committed to the existence of massive particles as the realist. According to the realist, we should treat the sentences of mature scientific theories as correspondence-true, and this treatment commits us to the ontologies of those theories. According to the NOAer, we should treat the sentences of mature scientific theories as deflationary-true, and this treatment also commits us to the ontologies of those theories. The difference in truth-theories is a difference over the nature of the truth-predicate, not a difference over the existence of the entities postulated by scientific theories. Insofar as interpretation of beliefs in scientific theories are concerned, then, NOA is not distinct from realism: a literal interpretation is a realist interpretation. [43]

In short, the difference between the theories of truth adopted by realism and NOA fails to yield a substantive difference of interpretative stances with regard to ontology. Indeed, Fine acknowledges that NOA's ontology is not different from realism's:

> NOA sanctions ordinary referential semantics and commits us, via truth, to the existence of the individuals, properties, relations, and processes, and so forth referred to by the scientific statements that we accept as true. Our belief in their existence will be just as strong (or weak) as our belief in the truth of the bit of science involved, and degrees of belief here, presumably, will be tutored by ordinary relations of confirmation and evidential support, subject to the usual scientific canons (1984: 98).[44]

If we adopt the semantic interpretation of NOA then NOA collapses into realism. So the semantic interpretation cannot be the one intended by Fine. This conclusion is bolstered by Fine's extensive discussions of essentialism. He says that realism and anti-realism see

[43] This is hardly an original criticism. In conversation, it is probably the most ready response one hears to Fine. In print, it is forcefully argued by Musgrave: "the NOA [sic] is a realist who avoids desk-thumping, foot-stamping, and shouting.. Whereas older realists shouted and stamped in opposition to anti-realist conceptions of truth, on NOA's Ark realists content themselves with a 'stubborn refusal to amplify' their referential semantic conception of truth" (48).

Musgrave goes on to advance a more sophisticated interpretation of NOA, according to which NOA rests on Putnamian scruples about correspondence truth and metaphysical realism. Musgrave cites Fine's appeal to the problem of access in criticizing correspondence truth and parlays this into an accusation of conceptual idealism (53-60). I don't mention it in the text because this isn't a terribly popular interpretation of Fine's views. Though one can find textual support for it (not least of which is footnote #1 in Fine's (1984) that expresses an affinity with Putnam and Rorty), it seems clear that Fine intends his position to be a bit more subtle than those views (at least, as far as philosophy of science goes). I also don't mention it because it is pretty obviously wrong: it requires that the NOAer remain agnostic about "whether or not [theoretical] statements are to be taken at face-value" and about "how they are to be interpreted [and] what their ontological commitments are" (60). But, as I show in the text, NOA is agnostic about none of these things: the ontological commitments of the NOAer are made on the basis of reasonable scientific theories, and those theories (including theoretical statements) are to be interpreted at face-value.

[44] See also (1986: 176), (1996a: 150).

science as having a definite aim (for realism, the truth; for anti-realism, instrumental reliability or empirical adequacy) and they interpret the sentences of science accordingly. NOA, on the other hand, sees no aim for science; it is content to "take things just as they come" (1996b: 253) and so refrains from interpretation. This suggests that the intended difference between realism and NOA lies in their differing conceptions of the scientific enterprise.

But what of Fine's explicit claim that truth is what separates realism and NOA? This interpretation of NOA -- let's call it the anti-essentialist interpretation -- can easily accommodate the claim. If realism is committed to the essentialist thesis that science aims to provide true claims about the world, then the realist is likely to adopt the correspondence theory of truth. After all, what better way to yield truths about the world than to look to the world as the truth-maker for one's sentences? NOA does not see science as aiming at the production of true claims and so has no reason to adopt this theory of truth. Hence the difference in truth theories. But the difference is not fundamental; it is a result of the deeper divide over the question of the aims and goals of science.

There is some support in the literature for this anti-essentialist interpretation of NOA. The most succinct statements of it are offered by Abela (1996) and by Brandon (1997). Abela tells us that "Fine opposes the idea that science requires a rational reconstruction of its practices. Science, we are told, creates its own story. It's goals, aims, and practices 'occur spontaneously and locally'" (73). According to Brandon, "Fine claims that [realism and anti-realism] assume an aim or essence for scientific activity, an aim that provides us with a benchmark for sorting the sheep from the goats. They presume that science can only be philosophically understood by reference to some such extra. NOA, on the other hand, just takes what it finds..." (233). Both characterizations contain crucial turns of phrase that are, at best, metaphorical ('occur spontaneously and locally', 'just takes what it finds'), and need to be cashed if we are to make a go of this anti-essentialist view of NOA. Nevertheless, the broad idea is clear: NOA resists the essentialist impulse to attach to science a pre-determined set of goals and aims, and this is what sets it apart from realism.

Martin Bunzl's anti-essentialist reading of Fine goes some way towards clarifying things. He contends that "[f]or Fine, both realists and antirealists err in their failure to appreciate the characteristics of science as a social enterprise" (1994: 450). On this view, realism and anti-realism "are to be understood as competing claims to be applied uniformly to science as a whole rather than as piecemeal theories about particular scientific theories" (451) while NOA, as Fine puts it, "thinks of science as an historical entity, growing and changing under various internal and external pressures. Such an entity can be usefully studied in a variety of ways- sociological, historical, economic, moral, and methodological- to name a few" (1986: 172). Realism and anti-realism fail to respect these several aspects of science, trying instead to force it into particular philosophic frameworks. NOA, by contrast, recognizes the fluidity of science as an enterprise and so adopts no single framework within which to analyze its theories.

Joseph Rouse also argues for a reading of NOA along anti-essentialist lines, contending that "an historically sensitive and open-ended particularism is ... fundamental to the natural ontological attitude" (1991: 625). By 'particularism', Rouse has in mind the contextualization to "actual scientific use" of "supposedly philosophical concepts" (612). The idea, I take it, is that claims to existence, explanation, etc., are to be evaluated not against the backdrop of a universal philosophic theory (such as realism or anti-realism), but in the context of the more localized concerns and constraints of the relevant field of scientific investigation.[45] Realism, we are told, is "global and essentialist" and "concerned with whether unobservable entities exist" while NOA is "local and pragmatic" and concerned with (e.g.) "whether gravitational lenses or releasing hormones exist" (611).

This last contrast states the intended difference rather nicely. Realism is concerned with broad philosophic questions and attempts to interpret the claims of scientific theories accordingly. NOA, by contrast, has no such broad concerns and so does not so interpret theories' claims. Hence realism takes an interpretative stance – required by its essentialism – while NOA takes none. Particular theories' existence claims are interpreted by the realist as confirming instances of realism, and particular theories' failures are interpreted by the anti-realist as confirmations of the pessimistic induction.[46] NOA makes no such demands of theories: if a given theory says it is reasonable to

---

[45] This talk of 'fields of scientific investigation' is Rouse's (611-612), and it's not clear whether he has in mind stages in the history of science or different kinds of science (i.e. physics vs. geology) or something altogether different. The talk of historical sensitivity would seem to indicate the former view, but the talk of actual scientific use (and the overall tenor of Rouse's comments) seems to favor the latter view.
[46] Anti-realists who do not endorse the pessimistic induction take the failures as confirming instances of the failure of realism, *sui generis*, I suppose.

believe in the existence of X's, then the NOAer believes X's exist, *punkt*, and draws no moral from this.

Seen in this light, the anti-essentialist interpretation of NOA comes to resemble the naturalistic interpretation offered by Leplin (1997: 173-177). Leplin casts NOA as a philosophical species of minimalism: it asks us to "rule out of court philosophical questions about science" (174) in favor of the view that "scientific theories 'speak for themselves'" (175). This looks like a close cousin of the view espoused by Rouse and Bunzl: NOA urges us to resist the urge to draw big conclusions from local scientific practices; instead, we ought simply to accept local scientific practice. Why should we do this? Because, as Abela and Brandon note, Fine thinks it is a mistake to attribute to science intrinsic aims or goals. There is nothing that all scientific practices and theories have in common – no striving after a common goal (the truth or empirical adequacy or anything else), no common methodology, no common essence. All there is, is local practice and particular theories.

What, on this reading of things, are we to make of the putative distinction between realism and NOA? If we take realism to consist in the essentialism of which Fine accuses it, then NOA and realism are trivially distinct: we are committed to the ontologies of particular theories, and so to realism about the theories of science, but not to realism about science generally. We must assume, I think, that Fine doesn't intend this: NOA is meant to be a substantive alternative to realism. But if we take realism to be about ontology, as Fine does, the distinction is in danger of collapsing.[47] If being a realist means that one is committed to the existence of those entities postulated by mature scientific theories, then it would seem that NOA is a species of realism- for NOA, too, is committed to the existence of those things that mature theories say exist.

Perhaps, however, Fine conceives of realism in a more complex way than merely as an essentialist program. Suppose he conceives of it as the thesis that scientific theories build upon one another, offering increasingly accurate accounts of the nature of things. Suppose Fine's realist thinks that there is good reason to be committed to the existence of those things that mature theories say exist *and* that new theories refine the commitments made by previous theories. In short, suppose Fine thinks that realism has both an ontological aspect and a cumulative aspect. Then the rejection of essentialism would be sufficient to make NOA non-realist, but insufficient to distinguish it from realism when it comes to ontology.

This is a reasonable conception of science- it certainly is not idiosyncratic, and it is in keeping with what Fine has to say in his several polemics on realism and anti-realism. If it is his view of realism, then we can account for the similarity between realism and NOA and at the same time make good on Fine's claim that the two are distinct. On this view of things, Fine's bone of contention isn't ontological; it's methodological. The concern isn't that unobservable entities exist—it's not about existence at all. Rather, it's about theory

[47] For reasons to think that Fine conceives of realism in ontological terms, see his (1991: 80), where he characterizes realism as having "a double aspect: epistemological-cum- metaphysical." (Be sure to see the footnote also, where he approvingly cites Devitt on this matter). See also (1986: 150), where realism is cast as a metaphysical thesis that has a semantic side to it. Fine (1984: 85) characterizes the argument for anti-

change. NOA objects to the view of science as a progression towards the truth.[48] Recall our introduction to NOA:

> As a scientist, say, within the context of the tradition in which he works, the NOAer, of course, will believe in the existence of those entities to which his theories refer. But should the tradition change, say, in the manner of the conceptual revolutions that Kuhn dubs 'paradigm shifts,' then nothing in NOA dictates that the change be assimilated as being progressive, that is, as a change where we learn more about *the same things* (1984: 98, emphasis in original).

By rejecting the essentialist claim that science strives after a single goal and is properly characterized by a single methodology, Fine is able to distance himself from any reasons to be a progressivist about science. Locally (to use his language) there is good reason to believe scientific theories. Globally, however, there is none: particular theories command our ontological assent or dissent, but science as a whole does not.

On this reading NOA is a hybrid position. In terms of ontological commitment, it is thoroughly realist. In terms of theory change, it is anti-cumulativist. Fine is perfectly comfortable with commitment to the existence of what scientific theories tell us exists, and this is a straightforward realism about ontology. Fine is also clear about NOA's resistance to the cumulativity thesis: he says that there is nothing in NOA that requires us to regard theory change as progressive. This anti-cumulativism is bolstered by the claim that there is no overarching goal or method of scientific inquiry: if there is no such common element, then there is no reason to think that two theories occurring in

---

realism as concluding in the claim that "the entities mentioned in explanatory principles [of scientific theories] need not exist."

[48] Schalgel (1991: 320) is the only commentator who seems to even acknowledge this aspect of Fine's position. Yet he gives it only a paragraph's attention, criticizing the view and moving on without stopping to consider the role of the view in the overall theory.

succession have anything in common. The anti-cumulativism is also bolstered by the rejection of the correspondence theory of truth: if there is no commitment to 'a way the world is' then there is no reason to think that successive theories offer better or worse accounts of the world.

Another way of characterizing NOA is as a version of weak epistemic realism. The epistemic realist is committed to the thesis that we have evidential reason for believing in particular theories. The epistemic realist does not think that these theories are true, but that that a careful sifting of the relevant evidence points towards these theories, and none of it points to competing theories (or not as much does). On this evidential basis, the epistemic realist concludes, we should accept these theories. There is nothing in this that requires the linking of presently preferred theories with the ontologies of past or future theories. Hence the epistemic realist can, as Fine does, embrace realism about ontology and at the same time be an anti-cummulativist.[49]

This is a peculiar position. On the face of things, realism about ontology requires – on pain of irrationality – realism about theory change. If you want to think that your favored scientific theory tells us what exists, you had better have an account of theory replacement. One natural account of this is in terms of approximation to the truth: the new theory and the old theory both get things kind of right, and the new theory is closer to being right. But this is to buy into cumulativism, and hence into a stronger realism than epistemic realism. Rational realism about ontology, in short, seems to require cumulativism. (At the very least, realism without cummulativity makes theory change a mystery.) That Fine tries to have it both ways is what makes NOA at once a confusing and frustrating doctrine: it is confusing because Fine doesn't make explicit that this is what he is trying to do, and it is frustrating because this is something that most people think can't be done.

Fine's trick is to go naturalistic. For Fine, there is no *a priori* answer to the question 'how does science work?' All we have to go on are the actual practices of working scientists. And working scientists do not worry about the pessimistic induction. They are concerned with their own research projects and research programs- with determining, as Rouse says, whether releasing hormones or gravitational lenses exist. If they think these things exist, then we should, too. If they don't worry about the place of those theories in the unfolding history of science, then we shouldn't either. When you think about it, it's sort of a naturalized version of Kuhn: all there are are paradigms and within paradigms there are no concerns over whether old theories had it right or new theories might prove current ones wrong. All there is is a commitment to current theories, and the ontological commitments and methodological practices that they involve.

Fine is arguing that the only legitimate thing to do when asked how theory replacement works is to beg out of the discussion: all we've got to go on are particular theories and particular scientific practices, and this simply doesn't determine an answer to the question of whether theory change is cumulative or not. Strictly speaking, then, Fine is not arguing for anti-cummulativity, but for agnosticism about theory replacement. Indeed, he says that "adherents to NOA are free to examine the facts in cases of paradigm

[49] Seen in this light, Fine's position is very similar to Laudan's.

shift, and to see whether or not a convincing case for stability of reference across paradigms can be made without superimposing on these facts a realist-progressivist superstructure" (1984: 98). The problem with realism, he thinks, is that it wraps progressivism in with the ontological aspects of realism, and forces one to buy the whole package. What Fine is trying to do is separate out the issues concerning theory change from the ontological issues. He doesn't think we have much of a choice with regard to the latter set of issues – we must believe that which science tells us to believe – but we do have a choice with regard to the former.[50]

3.1.2 NOA and success

Does Fine's naturalistic gambit work? There are, to be sure, a host of issues that can be engaged in developing an answer to this question. What reasons does Fine offer for thinking that agnosticism about theory replacement could be right? Can the combination of realism about ontology and anti-cumulativism be sustained? Does agnosticism about theory change weaken the irrealist element of NOA to such an extent that NOA risks again collapsing into realism? These are interesting issues, but they are not issues that would be usefully taken up in this inquiry. All we are after, recall, is a determination of whether NOA offers a distinctively non-realist account of the success of science. In order to make that determination, we needed to figure out exactly what NOA maintains, and that led us to an analysis of the several interpretations of NOA available in the literature. None of them yielded an account of that position that saved it from collapsing

[50] Brandon's criticism of NOA takes on a certain urgency when things are cast in this light. Brandon argues that Fine has not motivated his bias towards the social practice of science as opposed to, say, religion or witchcraft. But if we can choose whether or not to accept the claim that science is progressive, why can't we choose whether or not to accept the claim that science is revelatory?

into realism, and we advanced the hybrid account as an alternative to them, with the naturalistic gambit and agnosticism about theory progression a part of that account. At this point, NOA is a well-defined position that differs from realism, and we can get back to the business at hand: can NOA explain success? Like everything else having to do with NOA, the answer is 'yes and no.'

Fine's official line on success seems to be a dismissal of the success argument as a whole. Calling it 'the explanationist defense of realism,' he catalogs four possible responses to it (1991: 82-83), each of which finds some measure of support in Fine's varied writings on realism. He first considers rejecting outright the question of success, contending that "one might not accept the presupposition that the instrumental success of science needs (or 'requires') explanation." (82)[51] We'll see van Fraassen make a similar move in response to the realist, contending that not all demands for explanation are equal. This response might be buoyed by the argument in Fine (1986) that every realist explanation of success has an instrumentalist surrogate.[52] If this is the case, then the demand for explanation cannot yield a determinate answer – for each realist response, there is an equally good anti-realist one – and we would do better to reject the call for explanation.

[51] Notice that Fine is talking here about the instrumental success of science, or the methodological success of science. We'll look in a few pages at what Fine has to say about predictive success, but consider: insofar as Fine advocates a localized realism about particular theories, he can account for predictive success just as the realist does. But, whereas the account of predictive success underwrites the realist's account of methodological success, Fine's anti-cumulativism forces him to drive a wedge between predictive success and methodological success. Hence he needs to explain (or explain away) methodological success independently of his localized realism.

[52] The argument is simply that any evidence for the existence of an unobservable entity is also evidence for the instrumental reliability of the posit of that entity, and there are no non-question begging methods of establishing that the realist conclusion is better than the instrumentalist one. See Section 3, especially pp. 162-166.

He then suggests that "we can challenge the explanandum, that science is instrumentally successful" (82). This move, familiar from Laudan, exploits the large number of failures in the history of science and tries to minimize the wonder of occasional success. We'll consider this argument in some detail in the next chapter; suffice it for now to say that Fine endorses it in both his (1984) and his (1991) criticisms of the abductive argument for realism.

He also considers challenging "whether any explanationist defense of realism is reasonable in the context of a debate over the reliability of the hypothetical method" (82). This is the familiar accusation of question-begging: the legitimacy of abduction is precisely what is in question in the realism debate, and so abductive defenses of realism are suspect. Fine's metatheorem (1984: 84-87) is intended to establish precisely this claim: "to argue for realism one must employ methods more stringent than those in ordinary scientific practice. In particular, one must not beg the question as to the significance of explanatory hypotheses by assuming they carry truth as well as explanatory efficacy" (85-86).[53]

Finally, Fine argues that "in the contest with instrumentalism in particular, the explanationist defense of realism seems especially poor" (82-83). The reason is the same as that for rejecting the call for explanation: there is an instrumentalist transform of every realist explanation of success, namely, the replacement of instrumental reliability for each instance of 'truth' in the realist's argument. An explanation of instrumental success in terms of reliability will be just as effective as an explanation in terms of truth – will save all the same phenomena – and so will allow instrumentalism to compete on all fours with realism when it comes to explaining success.

Fine concludes the discussion and moves on without giving an indication of which response he favors. He seems to support all of them and to regard their collective force as sufficient to demolish the explanationist argument. We'll have opportunity to consider these responses in some detail in the next chapter, when we examine the realist explanation of success. Our aim now is to glean what sort of explanation of success we can expect from NOA. Insofar as Fine supports the first and second of these responses, we can't expect *any* explanation from him, for these indicate that Fine wants to opt out of the game. The third and fourth responses indicate that he thinks there is a real phenomena to be explained but that neither realism nor anti-realism is up to the task. If this is the case, then there should be some account of success on offer from NOA.

We get a hint of Fine's preferred response in his (1984: 100) discussion of the problem of the small handfuls. This problem, recall, is that scientists regularly consider only a small handful of the large number of possible theories appropriate to some area of inquiry, yet the theories they choose often succeed. NOA's account of this echoes the second response above: "the background [to the problem] was to keep in mind that most such narrow alternatives are not successful. I think that NOA has only this to say." Fine then

[53] In the previous chapter, we saw Richard Boyd advance a naturalistic response to this challenge, contending that his meta-abduction (from the success of the scientific method to realism) works as part of an 'overall realist package' to justify the abduction.

offers a curious sketch of a solution that centers on the preservation of reference as a basis for guessing before returning to the theme of frequent failure:

> If you believe guessing based on some truths is more likely to succeed plain and simple, then if our earlier theories were in large part true and if our refinements of them conserve the true parts, then guessing on this basis has some relative likelihood of success. I think this is a weak account, but then I think the phenomenon here does not allow for anything much stronger since, for the most part, such guesswork fails.

And that's it. All we hear from Fine on the matter of NOA's explanation of methodological success is that there really isn't much of a phenomenon to be accounted for and, insofar as there is, we can account for it by appeal to guesswork based on earlier theories that were in large part true. The latter part of this account is, as I say, curious, for it is at odds with what comes before and after it, and it is inconsistent with Fine's stance towards the preservation of reference and the cumulativity thesis.[54] Even if we allow Fine this appeal to reference, however, his position is that there is no problem of the small handfuls because theories fail more often than they succeed. NOA's party line on the issue of methodological success is to deny that there is any such success.

What about predictive success? Here Fine cannot be so dismissive, for he wants to allow that scientific theories can yield answers to questions like 'do releasing hormones exist?' And surely part of a scientist's answer to such a question is going to involve experimental set-ups that involve predictions. Recall that, for NOA, ontological commitment is made on the basis of "ordinary relations of confirmation and evidential support, subject to the usual scientific canons" (1984: 130). The usual scientific canons hold prediction in pretty high regard, and so NOA, if it is to be an adequate account of how science works, must include an account of how theories manage to yield accurate predictions. Fine doesn't address this issue, but it is not difficult to figure out how such an explanation would work: since NOA allows that the entities and processes postulated by scientific theories exist, it can account for the predictive success of a given theory just as realism does.

Notice, however, that this explanation is going to be much weaker than the realist's. Fine can explain why a theory yields a successful prediction by affirming the ontological commitments of the theory and parroting the realist's explanation. But when a new theory comes along to correct the predictively successful one, Fine's up a creek. If he admits that the correcting theory and the corrected theory share reference, then he has renounced his agnosticism about theory change (in this case, at least) and so explains success only by becoming a full-fledged realist. If he refuses to acknowledge a preservation of reference, then he's at a loss to explain the success of the rejected theory. Suppose some theory T enjoys the requisite sorts of success. For standard reasons we accept T's ontology. Fine can explain T's success by appealing to its ontology. But suppose that some new theory T* comes along. T and T* posit incompatible ontologies, and T* enjoys all of T's successes plus a few of its own. Fine's explanation of T's success is no longer viable, for T's ontology has been rejected by the same scientific community whose assent to it lent initial warrant to belief in its ontology. So Fine can explain predictive success, but only for the limited case of an un-corrected theory.

---

[54] This charge of curiosity does not depend on the hybrid reading of NOA: even if that reading fails, Fine is still staunchly opposed to the progressivism inherent in realism and so cannot so easily appeal to

The epistemic realist might respond by denying that previous theories were successful in the first place.[55] All we have to go on is current evidence, she might say, and current evidence points to current theories, and points away from past theories, and that's all there is to the matter. Notice, however, that this move runs counter to the pessimistic induction, which Fine embraces. The pessimistic induction works precisely because the history of science is a history of false but successful theories. It is the successes of those theories which stands in need of explanation, and the upshot of the pessimistic induction is that realism is not up to the explanatory task. So Fine cannot claim this as a reason to reject realism and at the same time deny that previous theories were successful: he must either deny that they were successful (in which case he loses one of the motivations for rejecting realism) or he must accept that they were successful (in which case he is unable to explain successive successes.)

Since he's got other arguments against realism, Fine would probably do best to give up the pessimistic induction. But isn't this question-begging? If we base our denial of a theory's success on the fact that current science leads us to reject that theory's ontology, then we wind up with realism: by definition, a theory can be successful only if its ontology exists. The thrust of the pessimistic induction is that we've got a catalog of theories that were, by all standard measures, successful – they yielded novel predictions and offered cogent explanations – but were false. The response being considered here is to simply deny that they were successful, and to ground this denial in the claim only true theories can be successful. This begs the question in favor of realism.

---

preservation of reference in solving the problem of the small handfuls.

[55] Thanks to Michael Levin for suggesting this response.

John Greenwood suggests that the success of false theories can be accounted for by the fact that false theories identify "systematic relations at the empirical level that hold for reasons unrelated to the reasons" postulated by the theories. Greenwood thinks such success would not be accidental or coincidental. I think it's clear that such success *would* be coincidental: the success of the theory would be due to a coincidence of facts that actually obtain and facts falsely postulated by the theory. Regardless, though, notice that an account of the success of progressive theories in terms of coincidence is going to court precisely the same issue with which we began: either the coincidences are related or they are not. If they are, then we have embraced cummulativity (of coincidence rather than reference); if they are not, then the progressive success of the theories remains a mystery.

So how does NOA fare with regard to the question of success? Not at all well. It offers no explanation at all of methodological success; rather, it deflects the question by denying the problem and, insofar as it admits the problem, it offers a realist explanation in terms of preservation of reference. As for predictive success, it admits the problem but offers a solution that piggybacks on NOA's realism about ontology. To make matters worse, that explanation is inferior to the realist's, for NOA is unable to account for predictive success in the context of the history of science. All it can do it explain why a theory, considered in isolation from theories that came before it or are to come after it, yields accurate predictions. In short, NOA poses no threat to realism: it is distinct from it, but it fails to yield a distinctive account of scientific success.[56]

3.2 van Fraassen's *constructive empiricism*

van Fraassen's anti-realism is, to be sure, stronger than Fine's. He claims that the aim of science is not truth, but empirical adequacy.[57] Features of theories other than empirical adequacy– explanatoriness, simplicity, predictive power, etc. – are merely pragmatic features of those theories. Such features carry no epistemic or metaphysical weight; all that matters is that the theory save the phenomena, and this can be accomplished by extending our ontological commitments to observable phenomena only. Hence the aims of science can be achieved without believing in the existence of unobservables entities and processes. Since it is *prima facie* reasonable to adopt the most metaphysically minimal position available that allows us to achieve our aims, van Fraassen concludes that the most reasonable philosophy of science available to us allows us refrain from commitment to unobservables.

Much has been written on constructive empiricism. This commentary tends to focus on the observable / unobservable distinction and on van Fraassen's theory of explanation. These are deserving topics, for they are central to the argument for constructive empiricism and each of them is highly contentious, but they are not my primary concern in the present discussion. My aim is to determine whether the constructive empiricist can

[56] The broader lesson learned is that realism without cumulativism cannot explain success

[57] Throughout the discussion of van Fraassen, I follow him in taking the proper analysis of truth to be the correspondence theory.

Implicit in this is a conception of scientific realism as the view that science aims at the truth. I don't want to get bogged down in another discussion of the proper conception of realism. Since our focus here is on the question of whether constructive empiricism can adequately account for success, I don't think we need to: van Fraassen's position allows us to refrain from commitment to the existence of unobservables, and that's enough to make it anti-realist (in my sense). All that remains is to determine whether it can explain success.

adequately account for the success of science, and I will deal with these topics only insofar as they connect to this theme.

So what does van Fraassen have to say about success? He first offers an evolutionary account of success, according to which science is an evolutionary process and the reason our theories are successful is that only successful theories survive. His pragmatic theory of explanation yields a second account of success, for it has the result that theories do not have explanatory power, *simpliciter*. If this is right, no philosophy of science can identify the features of a theory in virtue of which that theory is explanatorily successful, and so, one might argue, it is unfair to criticize constructive empiricism for not doing so. On a related note, van Fraassen argues that explanation must stop eventually, on pain of infinite regress, and there is no reason to not stop it at the level of observable regularities. This allows him to resist the demand for explanation altogether.

I don't think any of these responses are able to save constructive empiricism from the criticism that it is unable to adequately account for the success of science. What I propose to do is introduce constructive empiricism in more detail and then go through each of the three responses it yields to the realist's demand for an account of success, arguing as I go that none of them is able to insulate van Fraassen from criticism.

3.2.1 constructive empiricism

van Fraassen (1980: 12) offers a concise statement of the constructive empiricist position: "Science aims to give us theories which are empirically adequate; and acceptance of a theory involves as belief only that it is empirically adequate." A theory is empirically adequate "if what it says about the observable things and events in the world, is true- exactly if it 'saves the phenomena'."[58] Acceptance, we are told, is different than belief: to accept a theory is to accept it as empirically adequate, while to believe a theory is to believe it to be true (46, 49). So, according to constructive empiricism, science aims to give us theories that say true things about the observable world, and acceptance of a theory involves just the belief that it does say such true things.

What this amounts to is the view that theory acceptance involves nothing beyond adequately accounting for all the relevant observable phenomena. This in turn commits the theorist to belief in all the statements of the theory that are about observable phenomena. But it doesn't commit her to belief in those statements that are about unobservable phenomena. Of course, it doesn't preclude such beliefs. But it doesn't compel them, and van Fraassen thinks it is a rather simple epistemic principle to not stick one's neck out any farther than one has to.[59] The anti-realist upshot is clear: science aims not at the truth, but at empirical adequacy, and so acceptance of scientific theories does not require belief in the truth of the theoretical sentences of scientific theories. It follows that science does not compel belief in the unobservable entities postulated by its theories.

[58] All references to van Fraassen are to (1980) unless otherwise noted.

[59] Worrall (1984:69) notes that van Fraassen's only argument (if it can be called such) for this epistemic principle falls short. What van Fraassen says is that "it is not an epistemic principle that one might as well hang for a sheep as for a lamb" (72). Worrall notes that the point of this maxim is that "if a little and a lot are both available, and if the penalty is the same in either case, then one might as well go the whole hog and take the whole sheep." van Fraassen's thinking here must surely be that the penalty is not the same for constructive empiricism as for realism, for with realism there are more commitments and with commitments comes (epistemic and ontic) risk. But Worrall's broad point is well-made: there should be more discussion of this central epistemic principle.

But in actual scientific practice, there are factors that seem to influence theory acceptance other than truths about the observable world. Theories are accepted and rejected on the basis of mathematical elegance, simplicity, scope, completeness, and so on. van Fraassen argues that these are pragmatic virtues: "In so far as they go beyond consistency, empirical adequacy, and empirical strength, they do not concern the relation between the theory and the world, but rather the use and usefulness of the theory; they provide reasons to prefer the theory independently of questions of truth" (88). So a simple theory is no more likely to be true than a complicated theory; any reasons we have for preferring the simpler one are pragmatic reasons.[60]

What about explanatory power? Most people think that, *ceteris paribus*, the more explanatory of two theories is more likely to be true.[61] Indeed, the argument from explanatory success to scientific realism turns on precisely this intuition: the most likely account of the explanatory success of a given theory is the truth of that theory. If this were true, of course, the game would be lost for constructive empiricism: there would be some consideration beyond empirical adequacy that figures into theory acceptance. van Fraassen's solution is to argue that explanatory power is, like simplicity and elegance, a merely pragmatic virtue. If this is the case, then the fact that a theory is explanatorily potent might be a reason to prefer the theory, but it's not a reason to think it's true.

---

[60] There's a rather large literature on simplicity that van Fraassen ignores; he takes it as obvious that simplicity doesn't correlate with truth (p. 90). McMichael (1985) challenges him on this. Insofar as my aim in the text is simply to explore the implications of constructive empiricism for the success argument, I shall ignore such complications in van Fraassen's theory.

[61] So, too, for predictive power. For some reason, van Fraassen focuses solely on explanation. Given the arguments of the previous chapter, one wonders whether he has developed the strongest version of realism

Why think that explanatory power is merely pragmatic? Borrowing from Salmon (who borrows from Reichenbach), van Fraassen conceives of events as having causal histories. The causal net of an event is the series of causal processes that led up to it. Science describes that causal net, he thinks, and an explanation of why an event occurs consists in "an exhibition of salient factors in the part of the causal net formed by lines 'leading up to' that event. Those salient factors ... constitute (what are ordinarily called) the *causes(s)* of that event" (124). This is not the whole story of causal explanation, however, for the explanatory context requires that attention be focused on certain features of the causal net. Which features these are depends on the facts of the particular case at hand: a feature that in one context might be salient will in another context be irrelevant. In the context of an autopsy, for instance, a relevant factor might be 'blunt force applied to the head', while in the context of a police investigation, a relevant factor might be 'his wife found out he was cheating on her.' Both might be true, but each one would be considered irrelevant – and hence non-explanatory – in the wrong context.

Accordingly, we can think of a request for an explanation as a request for the identification of some particular factor in the causal net of the relevant phenomenon, and we can think of an explanation as the identification of the relevant causal factor. Which factor ought to be picked out depends on the context of inquiry: though any factor in an event's causal net will be present, only a few of them will be explanatory. It might be true that the man died because his wife found out he was cheating on her, for example, but it also might be true that he died because of a blow to the head. The context of

---

possible- explanatoriness might reasonably be seen as a pragmatic virtue of theories, but it seems unlikely that predictiveness is. This will come up again towards the end of the study of van Fraassen.

medicine calls for a medical explanation, though, and the context of police investigation calls for a motivational explanation, and what the detective considers a sound explanation will surely not satisfy the coroner.

van Fraassen concludes from this contrastive analysis of explanation that an explanation is not "a relationship like description: a relation between theory and fact. Really, it is a three-place relation, between theory, fact, and context" (156). Hence he arrives at his pragmatic account of explanation, according to which an explanation is an answer that "is evaluated, *vis-à-vis* a question, which is a request for information. But exactly what is requested, by means of the interrogative 'why is it the case that P?', differs from context to context" (156). Indeed, we are told, the contextual nature of explanation goes even deeper: "the background theory plus data relative to the question is evaluated, as arising or not arising, depends on the context. [sic!] And even what part of that background information is to be used to evaluate how good the answer is, *qua* answer to that question, is a contextually determined factor" (156).[62]

3.2.2 success and evolution

Given this position, how are we to account for the success of science? van Fraassen's official line is an evolutionary one: the reason current theories are successful is that "only the successful theories survive- the ones which *in fact* latched on to actual regularities in nature" (40). He illustrates this response by pointing out the futility of asking why the mouse runs from the cat. That's a bad question, he thinks: "species that do not cope with

[62] I have left out the details of the theory in order to ease the exposition. We'll have a chance to see exactly how this works in a few pages.

their enemies no longer exist. That is why there are only ones who do" (39). Similarly, theories that are not successful no longer exist, and that is why the ones that do exist are successful. Hence their success is not surprising and does not stand in need of explanation.

This reeks of sophistry. Our question is, 'why are current scientific theories successful?' van Fraassen's answer is, 'if they weren't successful, they wouldn't have survived; the fact of their existence explains their success.' But we can fairly put to van Fraassen the further question, 'why have these theories rather than some other theories survived?' The response to this question cannot appeal to their success, for that would be question-begging. And the response cannot appeal to any features of these theories other than those that make them empirically adequate, for this would violate the rules of van Fraassen's own game. The constructive empiricist, it seems, is forced to offer a brute success story: 'current theories are successful, this explains why we have them rather than some others, and that's all there is to it.'

Jarrett Leplin is onto essentially the same criticism. He notes that "Evolution deals statistically with populations; it is silent as to the individual case. It explains not why particular birds fly south, but why the attribute of flying south tends to be promoted over time among populations of birds." He then goes on to note that evolution, thus characterized, cannot explain success: "Similarly, if we ask of successful theories *why* they are successful, we need an answer that goes beyond an explanation of why science in general produces successful theories; we need an answer that appeals to attributes that discriminate among theories. Why does *this* theory work, while others equally the products of diligence and preferred methods fail?" (1997: 8, emphasis in original). What is wanted is an appeal to some feature of successful theories that accounts for their success, and van Fraassen's evolutionary strategy precludes such an appeal. It is thus found to be wanting.

3.2.3 success and the pragmatic theory of explanation

In a footnote to his discussion of evolution and success, van Fraassen acknowledges that the realist appears to have the upper hand.[63] The realist's answer to such a question as 'why does this individual bird fly south?' would involve a claim that evolutionary theory is true and that the truth of the theory is what explains the individual bird's migration. van Fraassen attempts to block this by appealing to the pragmatic theory of explanation. The pragmatic theory of explanation, recall, has it that the explanatory power of the theory is a function of the appropriateness of its response to a why-question given the presuppositions of the explanatory context. It thus makes no sense to speak about a theory's explanatory power, *simpliciter*: a theory will only have explanatory power if it satisfies the particular desires for information that are brought to bear by a particular request for explanation. Hence the same theory can be highly explanatory in one context, and not at all explanatory in another.

It follows that evolutionary theory explains the particular case at hand (some particular bird's southward migration) only given a certain context of explanation. If you change the explanatory context (say, by changing the presuppositions of inquiry) evolution fails to explain.[64] Hence the realist's apparent advantage is illusory: the fact that we can account for evolutionary theory's successful explanation of the southward migration of some particular bird is reason only to think the evolutionary account works in this context, and we must allow for the possibility of other accounts working in other contexts. Explanations must be indexed to contexts in order to be evaluated, and so there can be no eternal sense in which one theory is more explanatory than another.

Thus we arrive at van Fraassen's second response to the realist.[65] The realist argues that he, but not the constructive empiricist, can identify those features of successful theories in virtue of which they are successful: successful theories are theories whose postulated processes and entities really exist. If what counts as a good explanation varies from context to context, though, the comparative advantage of realism is illusory. The features of successful theories are as varied as the contexts of explanation: what makes a theory successful is the salience of its proffered explanations, not the reality of its posits. On this basis, we can attribute to van Fraassen the following response: 'of course constructive empiricism doesn't identify the features of successful theories that explain their successes- no such features exist. But constructive empiricism does advance a theory of explanation that allows us to make sense of the success of scientific theories and, insofar as it does, it is superior to realism.'

---

[63] Footnote #34, p. 40. The migration example employed in the next sentence comes from (Leplin 1997).
[64] van Fraassen's treatment of the asymmetries of explanation (111-134, esp. 131-134) is an exemplar of how to change the presuppositions of inquiry in order to force a change in the efficacy of proffered explanations.
[65] I'm doing reconstruction here: van Fraassen employs the pragmatic theory of explanation to establish explanatory power as a pragmatic virtue rather than an epistemic one; he doesn't present it as a possible response to the success argument for realism. I'm suggesting that this analysis yields more he realizes; it allows constructive empiricism to offer an account of explanatory success. The position I describe here is similar to the one I defend later on.

This is consistent with van Fraassen's (1985: 247) characterization of explanation as a virtue of informativeness. And recall his suggestion that the length of a tower's shadow can explain the height of the tower if we suppose that the tower is part of a giant sundial.[66] The theory succeeds in explaining the height of the tower so long as we switch to the sundial context; if we take the context to be as typically assumed, the theory fails to explain. With sufficient imagination, one could do this for any theory and for any explanation you like. The outlandishness of the imagined contexts is irrelevant, for the point is not that explanations like the sundial one are actually good, but that any theory can be explanatorily successful, and that we cannot determine whether or not a theory actually is successful without first determining the context.

It would appear, then, that the pragmatic theory of explanation yields a thoroughly anti-realist response to the realist. Explanations in terms of truth miss the point; what accounts for a particular theory's success in explanation is not the truth of the posits of the theory (as the realist claims), but the salience of the appeal to those posits given the broader context in which the demand for an explanation is made and answered. Moreover, this response sets up an anti-realist account of success: a theory is successful just in case it perform the informational job demanded by the explanatory context at hand.

This sounds well and good, but when we start looking at the details of the pragmatic theory of explanation, we seem to run into trouble. van Fraassen's theory has it that an explanation is an answer to a why-question. Any given why-question is a request for

---

[66] In Boyd, et. al. 1991: 324. Originally van Fraassen 1977.

information, and any such request implies the truth of the proposition about which information is being sought. Following van Fraassen's lead, let us call such a proposition the *topic* of the question. To use Leplin's evolutionary example, suppose the why-question is 'why do birds fly south in the winter?' The topic of the question is *birds fly south in the winter.* The question has a *contrast class*, which might include 'why do birds fly south rather than north?', 'why do birds fly south rather than walk?', 'why do birds fly south rather than elephants?', and so on. Lastly, there is what van Fraassen calls *a relevance relation*, "which determines what shall count as a possible explanatory factor" (142). A proposition is relevant to a why-question just in case it bears the relevance relation to the topic of the question. The proposition is an answer to a why-question – that is, an explanation – just in case it is true and claims that the topic, but not the other members of the contrast class, is true.[67]

What jumps out at us is the truth-restriction on explanation: in order to explain, a proposition *must be true* and bear the proper relation to the topic (and to the contrast class). It would appear that van Fraassen is treating explanation as factive, after all. And this would seem to run counter to what has come before: we cannot determine whether a theory offers a successful explanation merely by looking to the context; rather, we must look to the context *and* to the truth of the proposed explanation. But if only true theories can explain then false theories cannot explain, and the central realist belief in the explanatory power of truth remains.

[67] p. 143. The theory of why-questions is described on pp. 141-146.

van Fraassen's account of explanation includes a discussion of what counts as a good answer to a why-question (read: what counts as a good explanation), and it is worth looking here for clarification. Following van Fraassen (146), let us take K to be an accepted theory plus background information, and let us take A to be the proffered answer to some why-question Q whose topic is B and whose contrast class is X={B, C, ..., N}. We are after an evaluation of the answer *Because A*. van Fraassen says that there are three methods of evaluation. The first "concerns the evaluation of A itself, as acceptable or likely to be true," according to which we are to "rule out *Because A* altogether if K implies the denial of A; and otherwise we ask what probability K bestows upon A" (146, 147). The second method "concerns the extent to which A favours the topic B as against the other members of the contrast class" (146). According to this method, "K must imply B and imply the falsity of C..., N. However, *it is exactly the information that the topic is true, and the alternatives to it not true, which is irrelevant to how favourarble the answer is to the topic*. The evaluation uses only that part of the background information which constitutes the general theory about these phenomena..." (147, my emphasis). The third method "concerns the comparison of *Because A* with other possible answers to the same question..." (146): *Because A* is a good answer if the consideration of it along with the relevant background knowledge raises the probability of the topic while lowering the probability of the other members of the contrast class.

This allows us to resolve the tension in van Fraassen's theory in favor of the contextual element. According to the first method of evaluation, *Because A* is a good explanation if it is acceptable or likely to be true- if our background knowledge does not imply its denial. There's nothing here that requires the factivity of explanation: consistency is far too weak to do the job, and 'likely to be true' isn't much better: an explanation can be likely to be true (in the light of faulty background information) and at the same time be false. The second method of evaluation has to do with favoring: *Because A* is a good explanation of *B* if K implies B and also implies the falsity of the other members of the contrast class. Again, this is weaker than factivity: K can imply B and A can at the same time be false. (Indeed, van Fraassen makes explicit the fact that favouring is not factive!) The mere fact that the theory-plus-background knowledge leads us to think B is true and the alternatives to it are false is no reason to think that A is true, even if A is consistent with our background knowledge. According to the third method of evaluating explanations, *Because A* is a good explanation of B if consideration of A raises the probability of B (and lowers the probabilities of the other members of the contrast class) more than the background knowledge alone. As before, this is consistent with the falsity of A. If our background knowledge is faulty in such a way that consideration of a new false proposition A would make B seem more likely, then we will have satisfied van Fraassen's evaluative method with an *ex hypothesi* false explanation.

What we learn from the discussion of evaluation is that, for van Fraassen, it's not truth that makes an answer a good explanation; rather, it's likelihood of truth, and this likelihood is based on background knowledge. So can the pragmatic theory of explanation yield an account of the success of science that is as good as the realist's? It might do the job for explanatory success, but there remains the problem of predictive success. I noted above that van Fraassen is silent on this matter. At first glance, one might think this silence is justified by the forward-looking aspect of empirical adequacy. An empirically adequate theory, recall, is one that accounts for all the phenomena- past, present, and future. One might think that this means that an empirically adequate theory is one that gets all the future phenomena right, that is, one that gets the predictions right. Perhaps- but there is no explanatory mechanism postulated to explain how it is that the theory gets the predictions right. The realist's (imagined) original query was: 'what features of theories does constructive empiricism point to in accounting for the success of those theories?' van Fraassen's (imagined) response was that, for explanatory success, there are no such features to be identified, for explanatory success is a contextual affair. But this move is no help to van Fraassen here. What counts as a good explanation might depend on what information one already possesses, but what counts as a good prediction surely does not. If the theory says that some unexpected thing will happen, and if that unexpected thing does happen, then this is a good prediction. If the theory says that something will happen and it does not happen then it is a bad prediction. Whether the prediction is expected is, of course, relative to the information we posses. But once we grant that it is unexpected, we want an account of how the theorist knew to expect it. We want, that is, to identify some feature of the theory that accounts for its predictive success. The realist offers such an explanatory mechanism. and this is *prima facie* reason to prefer realism to constructive empiricism.

3.2.4 success and the demand for explanation

This leads us to van Fraassen's third response to the realist's demand for an explanation of the success of scientific theories: such a demand is ill-conceived for. not only do theories not explain, *simpliciter*, but explanation must stop somewhere, on pain of infinite regress, and it might as well stop at the observable level. Comparing the realist's demand for explanations to Aquinas' First Way, van Fraassen presents the realist's argument thus:

> Everything that is to be explained, is to be explained by something else. That some things are to be explained is evident, for the regularities in natural phenomena are obvious to the senses and surprising to the intellect. So we must either proceed to infinity, or arrive at something which explains, but is not itself, a regularity in the natural phenomena. However, in this case we cannot proceed to infinity (205-206).

The point here is surely to ridicule: just as Aquinas's argument for the existence of God can be refuted, with Hume, by postulating the universe itself (and not God) as the terminus of the explanatory chain, so we can postulate observable regularities as the terminus of explanation. If we can understand the universe only by assuming God's will to be its cause, Hume wondered, "how shall we understand God's will? And if we cannot understand God's will, why not stop with the universe, which we could not understand in the first place?" (van Fraassen 1980: 213). Similarly, if we can understand observable regularities only by assuming unobservable regularities to be their causes, how shall we understand those unobservable regularities? And if we cannot understand them, why not stop with the observable ones?

Suppose we grant van Fraassen the futility of the demand for explanation. Then the realist's challenge – 'how does constructive empiricism explain the abilities of theories to explain and predict phenomena?' – is a bad question. There are observable regularities in nature, and acceptable theories latch onto them, and that's all there is to be said about the matter: to demand an account of the regularities is to slip "back into the theological demand for explanation for its own sake" (Gutting 1980: 123).

This rings hollow: while van Fraassen is undoubtedly right that explanation must stop somewhere, it is too easy to say that it can stop anywhere. Surely there are constraints on the matter, and surely these constraints are not subject to negotiation. Suppose there is a landslide. We want to know why the landslide occurred- perhaps to satisfy our curiosity (whether scientific or salacious), perhaps to prevent another from occurring. We do not acquiesce in the fact of it's occurrence; rather, we assume that there is a cause and we seek it out. This insistence on a determinable cause is a result of our rules of scientific engagement, and we expect any party to the scientific game to honor it.[68] Indeed, even van Fraassen honors this rule: explanations are legitimate, he admits; explanations by appeal to unobservable entities are not.

But it is also a rule of scientific inquiry that explanation must proceed by appeal to the most basic elements available: one does not stop seeking explanations until one reaches bedrock. And what counts as bedrock is itself scientifically determined. If current science acknowledges biochemistry, then explanation cannot stop short of it.[69] Of course, if the legitimacy of biochemistry (e.g.) is itself the subject of scientific debate, then it is reasonable to draw the explanatory line before one reaches it. But if its

[68] Notice that we expect even amateur scientists (that is, the folk) to honor it. If Jones leaves the room suddenly and Smith asks why, we expect the others in the room to hazard guesses- perhaps based on Jones' behavior, or apparent mood, or what-have-you. Premature acquiescence in the unknown is suspicious.
[69] Of course, it can stop short if the explanation we are seeking is on another level: say, we are detectives – rather than cognitive psychologists – seeking the murderer's motive. For the purposes of this discussion, I am assuming that the relevant explanations are scientific. I'll consider objections to this assumption momentarily.

legitimacy is not the subject of debate then, absent compelling reasons, one is obligated to extend one's explanations to that level. To do otherwise is simply to practice bad science.

I take it as uncontentious that the scientific community accepts unobservable entities. Certainly it accepts as legitimate explanations that employ them. This means that, while explanation must stop somewhere, it cannot stop at the level of observable regularities. The rules of science do not allow it.

I imagine van Fraassen responding in one of two ways. He might contend that he offers compelling reasons to disallow such explanation, *viz.*, his argument against the need for such explanations (and, more broadly, his positive proposal for constructive empiricism). Second, he might contend that the only thing that makes scientific explanation scientific is its reliance on scientific premises,[70] and so there are no rules governing the stopping points of scientific explanations that do not govern the stopping points of non-scientific explanations. Neither response is adequate. The former begs the question, for the argument against the demand for explanation is a premise in the argument for constructive empiricism: if the demand is legitimate, and if constructive empiricism offers no account of success, then the success argument for realism goes through.[71] The latter response fails because there is nothing in the demand for explanation that requires

[70] "To call an explanation scientific is to say nothing about its form or the sort of information adduced, but only that the explanation draws on science to get this information (at least to some extent)..." (135).
[71] This is the case even if the argument against abduction goes through: van Fraassen's argument for constructive empiricism is (roughly) that we can account for all scientific practices without appealing to the existence of unobservable entities. One of these practices is explanation, and so he cannot include constructive empiricism in his argument for not allowing such appeals.

scientific explanations to be distinguished by their form or content. Rather, it requires that the context of explanation be scientific, just as van Fraassen says.

The demand for explanation of observable regularities is legitimate. This legitimacy extends to the demand for explanation of the predictive successes of theories, and the resources available to constructive empiricism are too impoverished to do the job. The evolutionary account of success originally offered by van Fraassen fails as well, and we are led to conclude that there is nothing in constructive empiricism up to the task of explaining the success of science.

4. REALISM AND SUCCESS

Neither Fine nor van Fraassen can explain the success of science; the negative argument for realism appears to go through. Now let's look at the positive argument for realism: the claim that only realism can explain success. There are arguments in the literature that purport to refute this claim. Laudan thinks the history of science is one of successful but false theories, and he thinks this shows that merely explaining success is no reason to regard a theory as true. van Fraassen offers several arguments designed to show that abduction is not a truth-conducive inference strategy. Much ink has been spilled attempting to defend realism against these attacks, and it is here that I propose to join the fray: I will defend the criticisms of realism advanced by Laudan and by van Fraassen.

Notice that showing that the positive argument for realism is unsound would put realism and anti-realism on all fours, forcing the realist to look elsewhere for support. If we were to show only the negative argument unsound, we would have two competing explanations, and further work would have to be done to show the comparative superiority of one or the other explanation. Given the counter-intuitiveness of anti-realism, the realist would retain the advantage. If, however, the proverbial rug is pulled out from under the realist, the anti-realist is free to develop an explanation, however counter-intuitive: as the realist says, better to have an explanation than no explanation at all.

4.1 Laudan's confutation

Laudan sketches a success argument for a view that he calls convergent epistemic realism (CER). CER is characterized by commitments to the approximate truth and genuine reference of scientific theories in the mature sciences, to the cummulativity of such theories, and to the claim "these theses constitute the best, if not the only, explanation for the success of science" (220).[72]

By extension of the commitment to approximate truth and reference, CER is committed to the claim that "there are substances in the world that correspond to the ontologies presumed by our best theories" (220). In addition, we are "to assume that a theory is successful so long as it has worked well, that is, so long as it has functioned in a variety of explanatory contexts, has led to confirmed predictions, and has been of broad explanatory scope" (222). It follows that any argument for or against CER is an argument for or against the success argument for realism, as we are treating it in this inquiry. If Laudan's critique of CER holds up, we will have shown the positive argument for realism to be unsound.

The argument for CER goes like this:

1. If scientific theories are approximately true, then they typically will be empirically successful.
2. If the central terms in scientific theories genuinely refer, then those theories generally will be empirically successful.
3. Scientific theories are empirically successful.

4. (Probably) theories are approximately true and their terms genuinely refer.[73]

Premise one postulates a connection between approximate truth and success. But the concept of approximate truth is too muddled to underwrite the realist's success argument: "On the best known account of what it means for a theory to be approximately true, it does *not* follow that an approximately true theory will be explanatorily successful" (229). He offers as an example the Popperian conception of verisimilitude. If a theory meets this condition – if it has more truth content than falsity content - it does not follow that any given class of the theory's consequences will be true. "Indeed," Laudan notes, "it is entirely conceivable that a theory might be approximately true in the indicated sense and yet be such that all its consequences tested thus far are *false*" (229). In short, the conditional *if a theory is approximately true then it will be successful* is no better than the account of approximate truth that accompanies it, and the realist has not specified an account that can ensure the truth of the conditional.

But suppose he has. Granting, for the sake of argument, the cogency of talk of approximate truth, Laudan argues that success does not betoken approximate truth. His argument is based on the assumption that "a realist would never want to say that a theory was approximately true if its central terms failed to refer" (230). Given this assumption, Laudan argues that the historical record includes a huge number of successful theories

[72] All references to Laudan are to (1984) unless otherwise indicated.
[73] The argument is on p. 220. Laudan sketches and responds (233-242) to a second argument for CER that I do not need to address in this inquiry. That argument purports to show that the realist's claims to genuine reference and approximate truth can be established by appeal to the 'limiting case' relations among successive scientific theories. Laudan's confutation of it takes the same form as his confutation of the argument described in the text. Since this is not the argument I am after in this inquiry, and since this is not the argument that Laudan's critics focus on, I ignore it here.

whose central terms did not genuinely refer and so which were not (by assumption) approximately true.

Premise two of the argument for CER is crucial, then. This premise postulates a connection between genuine reference and empirical success. Laudan rejects the conditional *if a theory's central terms genuinely refer then it will be successful* because, historically, there have been unsuccessful theories that contained genuinely referring terms (the chemical atomic theory of the 18^{th} century, for example, or Wegner's theory of continental drift). What about the converse? Laudan considers the conditional *if a theory is successful then its terms genuinely refer*. He rejects it because the history of science contains "a plethora of theories that were both successful and (so far as we can judge) nonreferential with respect to many of their central explanatory concepts" (231). He then offers a list of such theories which, he claims, "could be extended ad nauseam...".[74]

In short, theories can be successful without referring, and they can refer without being successful.[75] Since a theory must refer in order to be approximately true, it follows that theories can be successful without being approximately true and that they can be approximately true without being successful. So success is not an indicator of either

[74] I am begging some questions here by assuming, with Laudan, that false theories can explain. We'll consider objections and responses to Laudan's argument and, by extension, to this assumption, in a few pages; for now – for the sake of exposition – let's simply assume that false theories can explain.

[75] Some readers will object to this sweeping claim, contending that explanation is factive and so false theories cannot be explanatorily successful. I don't think these readers would take issue with the claim that a false theory can account for the phenomena, or that it can allow us to answer why-questions; their problem is with the claim that false theories can *explain*. I'm happy to give up the epithet 'explanatory success' and talk about a theory's ability to account for the phenomena: my interest is in realism, not theories of explanation. Of course, it would then follow that the realist's inference from explanation to realism is question-begging, for I am granting that explanations must be true. The realist would then have to argue that accounting for the evidence (answering the why-question) is reason to think the theory is true, and so we lose no dialectical ground. I'll come back to this in a few pages.

approximate truth or genuine reference; the success argument for realism fails. The concept of approximate truth is so impoverished that it cannot underwrite an analysis of empirical success and, more substantively, the history of science drives a wedge between success and realism.

This historical gambit also gives rise to the pessimistic, or historical induction: since successful scientific theories have been shown false by later science, we can inductively conclude that presently successful theories will be shown false by future science. It's worth noting that the historical gambit and the pessimistic induction are distinct: the latter requires the former, but one can endorse only the former.[76] This complicates things: we need to separate the responses to the historical gambit from the responses to the pessimistic induction. Since the aim of this inquiry is only to examine the success argument for realism (and not to defend anti-realism about current science), and since the historical gambit is necessary for the induction to get off the ground, I shall limit my discussion of the responses to Laudan to those that focus on the historical wedge he attempts to drive between success and realism. Arguments that seek to refute the pessimistic induction by focusing on those of its features having to do with induction, specifically, will be left aside.

4.2 responses to Laudan

Laudan argues that there is no clear connection between empirical success and truth/reference. One way the realist can respond to this is to exploit the 'maturity clause' in the arguments for realism. He might thus contend that the failed theories that Laudan appeals to weren't in fact mature and so cannot form the basis for rejecting realism about mature theories (which are the only ones worth being realists about). The idea is simply to cast the criteria for what counts as a mature science - and hence a legitimate indicator of truth and reference - so as to exclude successful but false theories.

Laudan anticipates this response (231-232). One problem with it, he thinks, is that it makes realism vacuous: mature sciences are defined as theories in which reference obtains. A second problem is that it is overly restrictive. The point of the success argument is to explain why science is successful. If some immature theories of science are successful, and if realism only explains why mature theories are successful, then the argument fails to do its job. Laudan thinks that there are some such immature, successful theories, and so he rejects the maturity defense as inadequate for the realist's broader needs.

This gives rise to a second manner of response: the realist can deny that past theories were successful.[77] Past theories were regarded as successful, this objection goes, but ultimately failed and so in fact never were successful. We simply *thought* they were successful- but we were wrong. Hence they pose no threat to the inference from success to realism.[78]

---

[76] Moreover, one needn't reject current science if one endorses the pessimistic induction; all one needs to do is reject inferences from the success of current science to the truth of current science. One might – as Laudan does – accept current science on other grounds.
[77] This response was urged on me in conversation and correspondence with Michael Levin.
[78] On such a view, not even the heliocentric theory of the universe is, strictly speaking, successful. All we are able to say is that the theory has been accepted as successful, that all indications are such that it will be

There is a certain rhetorical appeal to this: we would renege on the theory's claim to be true in precisely this manner ('oops, I thought it was true but I was wrong!'). What the realist needs for us to renege on, though, is the claim that the theory is successful, and it's not at all obvious that we would do so. When we say that theory is successful, we intend to point out its accuracy in prediction, its usefulness in explanation, its applicability in the technological realm, etc. If a theory enjoys such successes, and if that theory is later shown to be false (as Newton's theory was, or as aether theories were), we will not withdraw our assessments of the theory's predictive, explanatory, and technological utility.[79] We would, perhaps, marvel at the fact that the theory is false yet still managed to yield such accurate predictions and explanations. We might even wonder how to explain the success of such a theory.[80] But we wouldn't deny the success: correct predictions don't become incorrect because the theory that produced them turns out to be false.

What's going on here, I think, is a confusion between our assessments of theories and our assessments of the results of theories. If a currently favored theory turns out to be false, we would, as suggested, renege and say that we thought the theory was right, but we were wrong. But the explanations and predictions and technological achievements yielded by

successful, etc. Alternatively, we can say that the theory is successful and, if we are ever proved wrong, we will have to renege and say that we only *thought* it was successful.

[79] Michael Levin urges me to rethink the status of explanation in all this- at best, he says, a false theory is a potential explanation: in order for it to be an actual explanation, it must be true. Hence a false theory might be predictively successful or technologically successful, but it cannot be explanatorily successful, for false theories simply cannot explain. See note 76. I employ the phrase 'explanitorily successful' for ease of discourse, but I'm not invested in the theory of explanation.

[80] We might say that the *apparent* success of the theory is a phenomenon to be explained. More generally, we might say that the *apparent success of false theories* stands in need of explanation, and we might argue that realism is not up to the explanatory task. To say this would be say that false theories are not in fact successful, though, and thus to take a different approach to the problem than the one I take in this paper. Hence I do not adopt this strategy here, though it might, in the final analysis, be the better one.

the theory would stand on their own; they remain right, regardless of the status of the theory that produced them. Airplanes would continue to stay aloft, even if Bernoulli's law turned out to be false. Our assessments of theories are independent of our assessments of the results of theories.

'But how could airplanes fly if in fact Bernoulli's law is false?' the realist demands. Indeed, how can we explain this success? It is the realist's incredulity that gives us the success argument, and, in turn, Laudan's confutation of it. But the realist cannot respond to this confutation by reiterating his incredulity; that would be, as Laudan says, "a monumental case of begging the question" (242). The realist cannot fathom any explanation of a theory's success other than the truth of the theory, and so he denies the possibility of a successful yet false theory. But denial isn't sufficient; what is needed is an argument. Insofar as there isn't one here, the response is no good.

It would appear that we can't fix it so the troublesome theories of the past failed to succeed. Can we fix it so that they succeeded to refer? By exploiting the principle of charity introduced by Putnam (1984: 143), we can make "retroactive reference assignments" in order to make the sentences of past theories come out true. We do this by interpreting past theories so that they refer to current entities that we think exist. This way the success of the theories continues to be rooted in the existence of the entities and processes they postulate, even if they do not, strictly speaking, genuinely refer. "Surely," Putnam tells us, "the gene discussed in molecular biology is the gene (or rather 'factor')

Mendel *intended* to talk about; or it is certainly what he should have intended to talk about!" (143).

Hardin and Rosenberg (1982) base their defense of convergent realism on this retroactive reference approach. One possible move for the realist to make, they say, is to sever the connection between reference and approximate truth and claim that false theories were approximately true even though they failed to refer.[81] This can by done by re-assigning causal roles from referentially empty assignments to entities that are currently taken to exist. So "Mendelian genetics is still represented as an approximately true theory, even though its central terms can ... plausibly be said not to refer. The causal role Mendel accorded to genes is parceled out to other entities. In brief, this is done by showing that the diverse units of genetic functions ... work together to give a false impression of unity ... which Mendel mistakenly took to be phenotypes" (607). A second possible move for the realist is to make a reference assignment in the manner that Putnam suggests: "the changes in genetic theory in the last century can be presented in a manner which suggests that the central theoretical terms of Mendel's theory do successfully refer. They may be taken to refer to configurations of DNA and their polypeptide products, even though of course neither Mendel nor any other geneticist before the 1950's realized that it was these sorts of things to which the terms refer" (607).

Kitcher employs this response as well. Using Fresnel's analysis of diffraction as a case study, he considers two accounts of reference. On the first account, "Fresnel's references are explicitly fixed through the descriptions he gives ... in characterizing the wave theory of light and in introducing the wave equation. Tokens of 'light wave' ... thus have their references fixed by the description 'the oscillations of the molecules of the ether'. Since there is no ether, they fail to refer" (146). On the second account of reference, "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' ... genuinely refer to electromagnetic waves of high frequency" (146). Kitcher argues that Laudan's refutation of realism requires the first account of reference, but fails if we use the second account, and he argues that the second account has "a lot to be said for it"; indeed, that "Fresnel's success is based on his insights into the propagation of transverse waves, and his faulty beliefs about the constitution of those waves is irrelevant" (147).

Does this retroactive reference defense work? Maybe, but the realism that results is awfully cheap.[82] Such a 'retroactive realism' might be enough to preserve the letter of the realist position, but it doesn't do much for the spirit. For look what it licenses: we can say, for any successful theory - anywhere, anytime - that its success is due to whatever entities most resemble those implicated in the theory and which are said to exist by current science. This sort of blanket defense allows us to account for the successes of our theories, but it tells us nothing about the structure of the world. As Worrall (1996) notes,

[81] Recall that Laudan's argument against the success / approximate truth connection turns on the assumption that reference is necessary for approximate truth. Hardin and Rosenberg directly attack this assumption (606).

[82] As Arnold Koslow once said in another context, 'it's the K-Mart version.'

"the notion of one theoretical entity approximating another or of one causal mechanism being a limiting case of another is extremely vague and therefore enormously elastic. But if the notion is stretched too far, then the realist position surely becomes empty. If black 'approximates' white, if a particle 'approximates' a wave ... then no doubt the realist is right to be confident that future theories will be approximately like the ones we currently hold. This won't, however, be telling us very much" (155-6). Yet this is precisely what results: if the success of Mendel's theory can be attributed to the reference of his terms to something he didn't claim existed, Mendel hasn't said anything about the structure of the world.[83] Our revision of his theory (on the basis of current science) might be instructive, but we learn nothing from the original theory.

Psillos (1996) is not moved by such accusations of vacuity. He thinks that "It is not just that some current posit has taken up the place of an abandoned posit as the hypothetical cause of a set of phenomena." Rather, "The current posit is ascribed some (but surely not all) of the attributes ascribed to the abandoned putative entity, attributes in virtue of which it was thought to produce its effects" (312-313).[84] The idea here seems to be that we do not just replace the faulty ontology of the old theory with something useful from the current theory. Rather, we choose our replacements on the basis of the properties

---

[83] John Greenwood suggests that this is an unreasonable demand to make of the realist: given that the continuity of reference between the replacement theory and the replaced theory, we can't expect the replacement theory to tell us something about the world. But this misses the point of the reply to the realist: if the actual posits of the theory are irrelevant to the theory's success, realism becomes impoverished. This is the thrust of my talk of learning from theories and theories telling us something about the world: revisionist reference makes realism too easy and, in so doing, empties it of meaning.
[84] To be fair, he thinks that the problem arises only for cases in which the central terms of successful theories have been shown to be non-referring and that, for many successful theories, this isn't the case (312).

attributed to the original entities (though of course not all of them), and so the original ontology does tell us something about the world- it guides the choice of a replacement.

This doesn't make the strategy any more palatable, though; it just describes it. No one has suggested that retroactive reference assignments are made randomly; the accusation is that hindsight allows us to account for success without advancing our knowledge of what the world is like. The strategy amounts, in effect, to the following: 'any entity that is similar to those described by the theory and that is held to exist by contemporary science is responsible for the success of the theory.' Since this account of theoretic success changes with the ontology of 'contemporary science', all it tells us is that we can manipulate past theories to bring them in line with current theorizing. This is realism by fiat: we force successes into the mold of current science, and thereby 'prove' that success implies realism.

So the retroactive reference strategy fails because it doesn't tell us anything about the world and, as such, is a defense of a realism that really is not worth fighting for. But there is a version of it that might fare better, and which might preserve the not-unreasonable intuition that current science gets right what previous science got wrong. This version enjoins us not to treat theories as successful *in toto*. Rather, we should look at those parts of past theories that were successful and then look to see whether the entities and processes implicated in those successes are preserved by present science. It is not enough for Laudan to show that some theory was successful and employed a rejected ontology; he must show that the parts of the theory that were successful also essentially employed the rejected ontology. Jarrett Leplin offers such a defense, contending that "where past theories have met the standards imposed for warranting theoretical belief, their eventual failure is not total failure; those of their theoretical mechanisms implicated in achieving that warrant are recoverable from current theory" (1997: 145). He illustrates the strategy with an analysis of Fresnel's theory of diffraction, arguing that Fresnel's use of the ether "amounts to background interpretation, inessential to his use of his theory to obtain specific diffraction patterns" (146-147). Other parts of the theory did involve the ether essentially, and those parts do not stand up to empirical tests. So the claim that Fresnel's theory was successful but false is misleading: the successful parts were not false, and the false parts were not successful.

Kitcher (1993) makes a similar move. He distinguishes "two kinds of posits introduced within scientific practice, working posits (the putative referents of terms that occur in problem-solving schemata) and presuppositional posits (those entities that apparently have to exist if the instances of the schemata are to be true)" (149). The ether, he claims, was a presuppositional posit- "rarely employed in explanation or prediction, never subjected to empirical measurement ... yet seemingly required to exist if the claims about electromagnetic and light waves were to be true" (149). Like Leplin, Kitcher dissects Fresnel's theory and preserves those parts of it that meet the conditions imposed by later science. The theoretical mechanisms involved in those parts of the theories are working posits - needed to account for the successful parts of the theory - and the theoretical mechanisms involved in the rejected parts of the theory - those parts of the theory that fail the tests of later science - are the presuppositional posits.

On its face, this proposed defense offers a sophistication that the retroactive reference defense lacks. We do not simply recast the referents of the theoretical terms of past theories in currently accepted terms; rather, we examine those theories and see which of their terms would have to refer successfully for the theories to have been successful. At this point, the matter would seem to be an empirical one: once we have generated a list of such terms, we look at current science and see if the references were indeed genuine. If so, the strategy (and realism) is vindicated.

On closer analysis, though, it seems that this defense of realism turns on the retroactive reference strategy. By way of summing up his thoughts on the Fresnel case, Leplin makes a rather off-hand remark that favors revisionist reference: "Much of Fresnel's conception was right, *by our lights*, enough of it to account for the novel success of his theory" (148, my emphasis). More telling is the suggestion that "the success of reference does not require the referring expression to be true of its referent" (147), and the claim that "It was because it was light that [Fresnel's] ether-theoretic description referred to, and because light propagates as waves in the diffraction phenomena he was investigating, that his predictions were successful, although his theoretical account was erroneous" (147). Finally, we should note Leplin's suggestion that the distinction between referential and attributive uses of definite descriptions supports Kitcher's (and his own) position on the reference of theoretical terms (147).

Insofar as Kitcher's argument is identical to Leplin's, we can conclude that his argument will falter on similar grounds. The defense is based on the hunch that the empirical analysis of past successful theories will reveal that "the working posits" of those theories did all the work, while "the idle posits"[85] of those theories will prove to be- well, idle. But how are we to determine which posits are working ones and which are idle? The appeal to Fresnel as a case study makes it seem as though figuring this out is a simple matter of appealing to the historical record- after all, we've got Fresnel's memoirs to look to in determining what he had in mind. But things are seldom so clear-cut. (Indeed, we have seen that matters are not so clear-cut even in this case!)[86] How are we to determine whether some part of a theory relied on some or other posit? If we take the theories at face value, the posits in question will appear to be working posits; if they didn't, then they wouldn't have made it onto Laudan's list to begin with. If we employ a principle of charity, the approach reduces to the retroactive reference strategy: which posits will be deemed idle will be the ones that get abandoned by current science.[87] Moreover, if we determine retroactively whether a posit is idle then we haven't shown that idle posits aren't important to the success of the theory (and so concluded that they exist); rather, we have shown that they don't exist and so concluded that they aren't important to the theory's success. This doesn't show that the inference from success to realism is upheld by the historical record; it shows that, with hindsight, we can re-cast successful theories

---

[85] Kitcher 142.

[86] I don't want to quibble about how many case studies each side in the dispute has, but it is worth noting that neither Leplin nor Kitcher offers cases beyond the Fresnel case. Can their retroactive reference strategy account for other cases? How do we handle, e.g., the successes of Newtonian theory without absolute space and time?

[87] John Greenwood thinks I am being unfair here: the Leplin/Kitcher strategy is to "find the parts of Fresnel's theory essential to generating the successes of that theory, then ask if they can be interpreted as co-referential with the entities recognized by contemporary theories." But, Greenwood thinks, this does not constitute retroactive assignment of reference because " in principle, the essential components might be rejected by contemporary theories." This just seems wrong: if a contemporary theory accepts the essential posits of an historically successful theory, then the theory is accepted and there is no issue; if contemporary theory rejects the posits, then a co-reference assignment will be made. If the contemporary theorist rejects the posit but does not make a co-reference assignment, then he accepts the (apparent) success of a false

in currently acceptable terms. There remains the question of why the past theories were successful.[88]

4.3 van Fraassen's argument of the bad lots

The success argument takes the form of an abductive inference, or an inference to the best explanation. Perhaps the best example of abduction comes from van Fraassen: "I hear scratching in the wall, the patter of little feet at midnight, my cheese disappears- and I infer that a mouse has come to live with me" (1980: 19-20). As van Fraassen reasons from effect (scratching, missing cheese) to cause (presence of mouse), so the realist reasons from effect (the success of science) to cause (the existence of the entities and processes postulated by science). van Fraassen argues that this inference strategy is flawed.

The first argument van Fraassen (1989) offers can be called 'argument of the bad lot' (cf. Psillos 1996). When we apply the abductive inference strategy to a given theory, what we are doing is saying that it offers the best explanation of any of the theories we have considered. This is enough to justify our favoring it over its rivals. But it is not enough -- on its own – to warrant the conclusion that it is true. This conclusion requires an additional step- we must believe that the theories we have at our disposal are more likely to contain the truth than not. But the available evidence doesn't warrant this additional

---

theory. There are no other options- so, in the case where original reference fails, the options are either retroactive reference or denial of the success argument.

[88] Michael Levin notes that sometimes we will be able to identify idle posits on purely logical grounds: when a prediction makes no mention of a particular posit, even though the theory countenances the posit, we can say that the posit is idle with respect to that prediction. Most cases are not this simple, though: most cases at least appear to essentially involve terms that we take to be non-referring: these are the cases Leplin

belief, so the abductive inference strategy fails to serve the realist's needs. As van Fraassen puts it, "We can watch no contest of the theories we have painfully struggled to formulate with those no one has proposed. So our selection may well be the best of a bad lot" (1989: 143).

It's important to be clear about the upshot of this argument. van Fraassen isn't arguing that inference to the best explanation is irrational: the bad lots argument is perfectly consistent with the claim that, if one must choose a theory, one should choose the theory that offers the best explanation of the phenomena at hand. But this doesn't mean that the chosen theory is thereby likely to be true; that stronger conclusion could be reached only if we had independent reason to think that the lot of theories from which we have to choose contains a true theory. But there is no such argument to be had. So van Fraassen is allowing the rationality (as a tool of theory choice) of abduction, but he is denying that there is a connection between being the best explanation and being true.[89]

4.3.1 first response: privilege

van Fraassen considers a response, which he attributes to Boyd (1985) and Friedman (1979), that claims "privilege for our genius. Its idea is to glory in the belief that we are by nature predisposed to hit on the right range of hypotheses" (143). Evolutionary

---

and Kitcher are trying to handle. Levin's strategy, while it might work for some historical cases, leaves open a wide range of apparently successful but false theories.

[89] The connection between the argument of the bad lots and van Fraassen's epistemology is highlighted nicely here. The bad lots argument yields the view that abduction is not rationally compelling- that one is not irrational if one does not employ it (though it may well be rational to employ it). van Fraassen's epistemology is centered on justifications of changes in belief, not justifications of belief: any belief is justified so long as there are no rationally compelling reasons to not hold that belief. (cf. the distinction between English law and Prussian law.) So beliefs reached as the result of abductive inferences are

psychology lends itself to this line of response: given the fact of our survival as a species, it must be the case that we have developed the capacity to light upon true theories about the world. But, as van Fraassen points out, "Our new theories cannot be more likely to be true, merely given that we were the ones to think of them and we have characteristics selected for in the past, because the success at issue is success in the future" (143). Species evolve the way they do because individual members of that species without a particular trait die, leaving only those that do possess that trait to reproduce. There is no trait 'lighting on true theories' that the human species has evolved to possess, though; at most, we have evolved to respond to our environment in certain ways (to duck when things are thrown at us, avoid extreme temperatures, etc.). But such responses are too weak to underwrite the realist's appeal to privilege; there is no reason to think that our evolutionary dispositions to respond to our environment also incline us to choose true theories.[90]

Psillos (1996) thinks there is another sort of privilege available to the realist. He argues that the scientist's background knowledge narrows "the space for hypotheses that provide a potential explanation of the evidence at hand" (39). The fact that a theory is considered by a scientist, given the wealth of knowledge that the scientist already possesses, provides reason enough to think it is true. In fact, in most cases there will not be a

---

justified – realism is not irrational – but not epistemically mandatory- there are no compelling reasons to not hold the beliefs, because the constructive empiricist alternative 'as-if' is available to us.

[90] van Fraassen also considers (144) a rationalist defense of the appeal to privilege, according to which we are likely to hit on true theories because this is most consistent with our other beliefs about god (that we are made in his image, etc.). Even if we were so created, though, an additional argument would be needed for the claim that we were created to develop true theories about things like quarks and laws of nature.

Doesn't this response also beg the question against the critic of abduction? The realist argues from effect (survival of the species) to cause (ability to light on true theories). The anti-realist can say that decision to be made: given the phenomena to be explained, the background knowledge will lead the scientist to a single theory, which we will deem true. It is only "when the background knowledge does not suggest just one theoretical hypothesis [that] explanatory considerations, which are part and parcel of scientific practice, are called forth to select the best among the hypotheses which entail the evidence" (39). In short, explanatory efficacy acts only as a sort of tie-breaker; in most cases, the theory's truth will be apparent from the very fact of its consideration.[91]

Our theories are privileged, then, by the fact that we adopt them on the basis of the evidence and previous scientific knowledge. Hence they are likely to be true, and not merely the best of a bad lot. van Fraassen is apt to respond to Psillos by questioning the assumption that adoption on the basis of evidence and previous knowledge betokens truth; Psillos's response is that this reply "rest[s] on a dubious … assumption, *viz.*, that evidence can never guide scientists to form approximately true theoretical beliefs. Even though evidence does not entail theoretical beliefs, it can support some theoretical beliefs up to a high degree, so that it would be unlikely that the beliefs are outright false and the evidence what it is" (40).[92] This doesn't answer van Fraassen's question, though. To be sure, the evidence and our background knowledge give us reason to adopt some theories and not others. And there is every reason to think that the adoption of particular theories on such grounds is rational. But what reason is there to think those theories are *true*? The evidence only lends support to one theory relative to the other theories under

---

this is an acceptable way to account for our continued existence as a species, but there is no reason to think that it is true, for such would require the additional assurances that are in question.

[91] This account resembles the account of success I will sketch later on.

consideration, and so it lends a relative support only: this theory is more likely true than the other theories under consideration, but there is no reason to think that it is true unless we assume that the batch of theories available to us is likely to contain true theories. And nothing at our disposal gives us reason to think this.

Psillos will respond: our background knowledge does the job- given what we know about the world, we have selected these theories as likely contenders, and the evidence at hand allows us to further narrow the search to this one theory. So we do have reason to think these theories contain a true theory: the reason is that these theories were selected on the basis of a wealth of knowledge of how the world works. But, van Fraassen will reply, there's a sort of selective adoption of theories at work here. Given the background knowledge, it makes perfect sense that the set of contender theories will be limited, and it makes equal sense that the evidence will operate to select one of these theories. But why should we think that the background knowledge leads us to true theories? Why not think instead that the background knowledge leads us only to empirically adequate theories? The realist will point to the success of our theories and contend that this is reason enough: 'if our background knowledge didn't lead us to adopt true theories, then how can we possibly account for the fact that our theories are successful?' But this, of course, begs the question; we are asking whether abduction is a truth-conducive inference strategy

[92] Psillos seems to be hinting here at a Bayesian argument for privilege. (Michael Levin also suggests a Bayesian argument for the required principle of privilege.) The response sketched in the text will serve as a response to such an argument.

because we want to assess the success argument for realism, so that argument cannot be invoked in a defense of abduction.[93]

Psillos retrenches by claiming that constructive empiricism, too, must beg some questions. As the realist has to show that the batch of theories from which the scientist chooses is destined to contain the truth, so does the constructive empiricist have to show that the batch of theories from which the scientist chooses contains the "really empirically adequate theory" (37). But this requires the same sort of privilege of principle that the realist inference to truth contains, and so constructive empiricism is in no better shape than realism. Ladyman, et. al. (1997) in their response to Psillos, argue that the conception of empirical adequacy at work here is inadequate- that any theory that saves the phenomena is an empirically adequate theory. Since there is no question that there can be more than one such theory, Psillos' talk of 'really empirically adequate theories' is off the mark.

Might this be too quick? Psillos is arguing that empirical adequacy has a forward-looking aspect: for a theory to be empirically adequate, it must save all the phenomena, including unobserved yet observable phenomena. But none of the theories available for consideration in theory choice can be said to empirically adequate, thus understood, without making an inductive leap from the observed to the unobserved (but observable).

[93] The realist, of course, is not persuaded: given a clearly successful theory, he will argue (with Michael Levin) that "maybe the *a priori* odds of us making up a true theory are small, but given the evidence for the theory, we can throw *a priori* worries out the window." But this re-statement of the success argument simply cannot constitute an argument for it: the success of science is not at issue here; what is at issue is the account of the success of science. van Fraassen argues that abduction is not truth conducive and that this poses problems for the realist's account. The realist needs to deal with this argument on its own terms; re-asserting the realist intuition gets us nowhere.

So "in order to claim that the best currently available theory is empirically adequate, an ampliative claim is needed , asserting that scientists have already hit upon an empirically adequate theory" (Psillos 1996: 41).[94]

Psillos is on to something here, but it's not what he thinks. He has shown that van Fraassen needs induction to get his program off the ground, not abduction. To claim that a theory is empirically adequate, given the forward-looking aspect of empirical adequacy, one must show that the theory will remain empirically adequate. But abduction isn't required for this job; simple induction will do. In order to do some damage to van Fraassen, then, Psillos has to show that he is not entitled to induction.[95]

Ladyman, et. al. (1997: 370) respond to Psillos along the lines predicted by Psillos (1996: 42): constructive empiricism needs some privilege to get off the ground, but not as much as the realist. Since the realist's ultimate goal is truth, he must somehow ensure that the background theories informing our theory choice are indicators of the truth. The constructive empiricist, by contrast, seeks only to show that our chosen theories are empirically adequate, and so requires a much weaker principle of privilege. The question now becomes, 'which principle of privilege do we strive after?' This brings us full circle to van Fraassen's epistemology: ought we to hang for a sheep or only for a lamb? van Fraassen thinks that there is nothing to be gained by striving after truth; the realist thinks there is much to be gained and so thinks the extra risk worth it. van Fraassen insists that

[94] Now is as good a time as any to remind the reader that the thrust of this chapter is not to save constructive empiricism, but to bury realism. Hence *tu quoque* defenses of realism like this one are irrelevant to the overall purposes of this work. See Ladyman, et. al. (1997: 308-309) for a similar reminder.

there is no non-question-begging reason to think that truth matters to scientific inquiry; the realist thinks there is. But notice that the discussion of the argument of the bad lots started with an attempt to establish the relevance of truth, and look where we have ended up. The realist might be right – truth might be vital – but abductive proofs of the matter aren't going to help establish this vitality.

4.3.2 second response: force majeure

A second move van Fraassen considers on the realist's behalf is what he calls the force majeure reaction. The idea here is to "try and provide arguments to the effect that we must choose among the historically given hypotheses. To guide this task is the rule of right reason. In other words, it is not because we have special beliefs (such as that it will be a good thing to choose from a certain batch of hypotheses), but because we must choose from the batch, that we make the choice" (144). van Fraassen argues that this move is doomed to fail: "Circumstances may force us to act on the best alternative open to us. They cannot force us to believe that it is, ipso facto, a good alternative" (144-145). He then goes on to consider scenarios in which a person is forced to choose among unfavorable alternatives (jump the crevasse or spend the night in the freezing cold; open one of N doors knowing that there is a 1/N chance of a tiger being on the other side), concluding that the forced decision to choose does not mean that the chooser regards any of the options (including the one chosen) as a good alternative.

[95] I am indebted to Michael Levin for pointing this out to me, though I don't want to suggest that Levin offered the suggestion in support of van Fraassen.

This is an odd discussion. Certainly van Fraassen is right when he says that we often choose the least unfavorable alternative available to us- not because we think it is a good choice, but because it is the best of a bad lot, and a choice is required. What the realist needs to show is that there is some reason to think that the theories we have at our disposal are true.[96] The force majeure move, even if works, shows only that we think the theories we choose (from the batch of those we have to choose from) are likely to be true.[97] The anti-realist can happily admit that scientists think this, for thinking it doesn't make it so. Scientists (and laypersons) may choose one of a range of options because they think it is the right one. But this is irrelevant to the question of whether they are in fact right, and so the force majeure move fails to do the job the realist needs it to.

4.4 van Fraassen's argument from indifference

The third move van Fraassen considers on the realist's behalf is to retrench from the position of abduction to that of probablism: to claim that "Despite its name, it is not the rule to infer the truth of the best available explanation. That is only a code for the real rule, which is to allocate our personal probabilities with due respect to explanation" (146). van Fraassen's rejection of this move, which he calls the 'argument from indifference', is quick: there are any number of theories, formulated and unformulated, that are in accordance with the phenomena to be accounted for. Most of these theories are sure to be false. We "know nothing about our best explanation, relevant to its truth value, except that it belongs to this class. So [we] must treat it as a random member of

[96] The realist, of course, thinks that there are such reasons. Perhaps there are. But the point to be made here is not about reasons to believe, but about reasons to choose. Force majeure suggests that we must choose a theory and that this somehow means that the chosen theory is true; the present discussion aims to show only that this response is wrong. The realist would be better off offering reasons to believe.

this class, most of which is false. [sic] Hence it must seem very improbable to [us] that it is true" (146). Without some independent reason to think the theories at our disposal are more likely to be true, we must treat them as members of the class of explanations, most of which are false. It follows that the theory we select is likely to be false.[98]

Psillos (1996: 43) argues that we do know more about our favored theories than that they are members of the class of potentially true theories. We know that they have passed all manner of tests and that they have qualified as the best available explanations of the phenomena. Hence there is no reason to treat them as random members of the set of contender theories. Psillos goes on to anticipate van Fraassen's reply, viz., that there are any number of unborn theories that are also contenders for the mantle of truth and, since there is no way of comparing the chosen theory with them, there is no way to claim that the chosen theory is more likely to be true. Hence the chosen theory is rightfully treated as a random member of this class of theories, and the problem re-asserts itself. But, Psillos retorts, this presupposes the truth of the thesis of empirical equivalence. Moreover, even if this thesis is true (and it is not clear that it is), the proper epistemic response is one of temporary suspension of belief- we should wait until we have evidence that decides between the options. van Fraassen would have us permanently suspend choice, but this is a position that needs to be defended and not merely asserted.

[97] I am assuming that van Fraassen's talk of 'good alternatives' is to be interpreted in terms of truth. If he has something else in mind, the argument doesn't even get this far off the ground.

[98] van Fraassen considers another form of retrenchment, according to which "the notion of rationality itself requires these features to function as relevant factors in the rules for rational response to the evidence" (146). He gives a Dutch Book argument against it (pp. 166-169).

Ladyman, et. al. (1997: 309-311) argue that the argument from indifference doesn't need to rely on the thesis of empirical equivalence. Rather than take the argument to show that the chosen theory is likely to be false, the argument from indifference can be taken to show that abduction requires an additional premise to establish the truth of chosen theories. The abductive argument moves from the explanatoriness of the chosen theory to the truth of that theory on the assumption that explanations, when they are available, are to be preferred to mysteries. But an additional premise is needed to establish the truth of the chosen theory, viz., that "there is (almost) always a unique best explanation." (309). Ladyman, et. al., suggest that the argument from indifference can be taken to show the need for this additional premise and to suggest that there is no justification for it.[99]

An odd argument, I think. For starters, there is little in the argument from indifference to suggest that it is about the need for this additional premise. Moreover, without some argument for the claim that there can be no ordering of explanations, the realist surely has the advantage: the best explanation is the one that explains the most phenomena with the simplest ontology. No doubt, we can argue about what counts as an explanation, and about what counts as a simple ontology, but these are issues the constructive empiricist can (and, in the case of explanation, does) take up separately. As far as the argument from indifference goes, it would seem that van Fraassen does need the thesis of empirical equivalence to make the argument go through, and the argument does not speak to the

[99] This retort to the realist would be bolstered by an appeal to van Fraassen's pragmatic theory of explanation. Psillos (44-45) notices this, contending that the move wouldn't work because explanatoriness is only the first step in theory confirmation. But this seems to miss the point: if explanatoriness is a

need for an additional premise establishing an upper limit to the evaluation of explanations.

Ladyman, et. al., have a second reply to Psillos. They say that it doesn't matter whether the chosen theory is actually, or only potentially, a random member of a class of equally good explanations. The mere fact that there may be such rivals is enough to make the abductive inference to the truth of the theory unwarranted. In effect, they're saying that because the theory might be false, the realist cannot justifiably claim that it is true. But surely this is too strong. Any theory might be false, and any claim might be false, but this possibility is not enough to make us all refrain from making claims to the truth. Rather, we stake our claims with an implicit proviso that we might be wrong: 'I am currently working at my computer (aside: unless I am merely dreaming I am).' What the constructive empiricist needs to do is make the possibility of the chosen theory's being false a live possibility. And this seems to require a defense of the thesis of empirical equivalence, as Psillos contends.

Which leads us to the third response to Psillos: "we are not dealing with a mere possibility here. Fundamental physics provides us with well known examples of empirically equivalent theories" (310). But this is too quick; as Ladyman, et. al., acknowledge, it is entirely possible that "the occurrence of empirically equivalent rivals in physics may well be quite exceptional, because of some highly peculiar features of physics itself" (310). Besides, even if there are genuine cases of empirical equivalence in

---

pragmatic virtue, as van Fraassen contends, there can be no unique best explanation. Rather, context will determine the ranking of explanations and the rankings will vary from case to case.

physics, this does not settle the matter in favor of constructive empiricism. For recall Psillos' response to van Fraassen on this matter: given a case of empirical equivalence, the proper epistemic attitude is one of temporary suspension, until such time as we can decide among the theories. The attitude constructive empiricism recommends is permanent suspension, and this is too strong.

Psillos also responds to the argument from indifference with a *tu quoque*: as the realist must treat his chosen theory as a random member of a class of potentially true theories, so must the constructive empiricist treat his favored theory as a random member of the class of potentially empirically adequate theories. The response to this should come as no surprise: the constructive empiricist allows that there can be indefinitely many empirically adequate theories, and he doesn't want to argue that one of them is better than the others. We encountered this exchange earlier, and we concluded that Psillos' argument turned on the need for the constructive empiricist to establish, non-ampliatively, that the batch of theories from which the scientist chooses his theories contains a truly empirically adequate theory, that is, a theory that saves all unobserved but unobservable phenomena. Since van Fraassen does not seem to have established this claim, the *tu quoque* worked. But not here: what makes the argument from indifference run is the possibility of theories equally as good as the chosen theory, and this is a possibility constructive empiricism embraces.

So the argument from indifference fails because it requires a full blown defense of the thesis of empirical equivalence if it is to get off the ground. The defense of the argument of the bad lots goes through, though: the realist seems unable to offer the additional premise justifying the ampliative inference from a favored theory's explanatoriness to the truth of that theory. Of course, this is enough for van Fraassen: if any of his arguments against abduction succeeds, then realism fails. Similarly, it is enough for our purposes: if abduction is impugned even once, then the success argument for realism is stopped.

4.5 abduction and scepticism

Thus far, we have considered and rejected defenses of abduction that try to establish the privilege van Fraassen claims is needed for abduction to succeed. Does the realist have any additional defenses at his disposal? He might take the offensive, arguing against the rejection of abduction. Either van Fraassen must reject all abduction or he must reject abduction to the existence of unobservable entities only. The former position would seem to result immediately in a *reductio*- we use abduction in our daily lives, and to reject it as unsound would be to question inferences as innocuous as the inference from the scratching in the walls and the disappearing cheese to the presence of a mouse in the house. But the latter position requires the defense of an ontological distinction between observable and unobservable entities.[100] It seems that, in order to reject abduction, van Fraassen must embrace one of two untenable positions, and this serves as a *reductio* of the rejection of abduction.

---

[100] At first, one might think that the agnosticism urged by constructive empiricism requires only an epistemological distinction- van Fraassen isn't claiming that the entities don't exist; only that there is no compelling reason to think they do. But in order to establish a class of things in which there is no compelling reason to believe, van Fraassen must establish the ontological distinction. (Thanks to Michael Levin for setting me straight here.)

On the face of things, van Fraassen's position is the latter one: he accepts what Psillos (1996: 36) calls horizontal abduction – abduction to the existence of unobserved but observable entities – and rejects what Psillos calls vertical abduction- abduction to the existence of unobservable entities. He defends the observable / unobservable distinction with his well-known anthropomorphic account of observability: "X is observable if there are circumstances which are such that, if X is present to us under those circumstances, then we would observe it" (1980: 16) and argues that, while the distinction between observable and unobservable may well shift (the epistemic community may widen given "significant encounters with dolphins, extraterrestrials, or the products of our own genetic engineering" (1985: 256)), the distinction is real.

But this isn't enough. After establishing the distinction, van Fraassen must defend his selective skepticism about the unobservable. The defense is outlined by Ladyman, et. al. (1997: 316): "Even supposing that in everyday life we routinely use IBE [inference to the best explanation] to go beyond the observed phenomena, we do not routinely introduce new ontological commitments. In the earlier case [of the mouse in the wainscoting] we already believe that mice exist, that is, we use IBE to conclude new facts about tokens of types that are already included within our ontological commitments." What is at stake in the realism debate, they think, is the use of abduction to introduce new kinds of entities- hence horizontal abduction is unproblematic, but vertical abduction begs the question against the anti-realist.

Psillos takes issues with this analysis. He argues that horizontal abduction does involve the introduction of new types of entities- for example, "positing an extinct type of animal both is an instance of IBE and does introduce new 'ontological commitments'" (1997: 371). Moreover, he claims, vertical abduction "does involve introduction of new instances of known types, e.g., instances of the virus HIV" (371). But these responses are not indicative of the standard uses of abduction. Indeed, the terminology, introduced by Psillos, is misleading. We can recast things in terms of ampliative inferences to known types of entities (call it familiar IBE) and ampliative inference to unknown types of entities (call it unfamiliar IBE). van Fraassen's claim is that unfamiliar IBE is suspicious, but familiar IBE is not, and Psillos' objection becomes irrelevant. Of course, Psillos can still ask what justification there is for such a selective skepticism, but van Fraassen has a ready answer to this: his empiricism requires that ontological commitments be made on the basis of direct evidence, and unfamiliar ampliation fails to provide that evidence.

All of this would seem to be irrelevant, though, for van Fraassen wants to take the first horn of the dilemma and reject all abduction. Ladyman, et. al. (1997: 311-314) indicate that van Fraassen thinks "there is no discrimination to be made between horizontal and vertical IBE ... van Fraassen's arguments are directed against IBE understood as a rule of inference, not as an inferential practice" (312). The idea here is that there is nothing irrational about abduction, but there is nothing rationally compelling about it, either. Viewed in this light, the example of the mouse in the wainscoting is an explanatory foil, not an endorsement of horizontal abduction: for any case in which it appears that abduction is required to explain our understanding of the world, we can employ an alternative rule – inference to the empirical adequacy of the best explanation – that does not require us to make the ampilative inference required by realism. Either inferential rule allows us to explain the disappearing cheese, and hence there is no rationally compelling reason to endorse one or the other rule of inference: the example "cannot provide telling evidence between the rival hypotheses" (van Fraassen 1980: 21). Hence there is nothing compelling about IBE – either the vertical or the horizontal variety – and so we are rationally free to reject it.

Of course, the fact that we are free to reject abduction isn't enough to justify van Fraassen's rejection of abduction. But if abduction isn't needed to make sense of practice then there are good methodological grounds for rejecting it. The realist thinks that the extra risk (of positing causes that do not exist) involved in ampliation is "a necessary consequence of aspiring to push back the frontiers of ignorance and to get to know more things, in particular about unobservable causes of the phenomena" (Psillos 1996: 42). But this extra risk is justified only if we agree with the realist that there are unobservable causes of the phenomena. If we disagree with this, then the realist's ampliation yields only ontological danglers and bloated ontologies.

But if we replace abductive inferences to the existence of causes with abductive inferences to the empirical adequacy of causal explanations, don't we thereby license skepticism about the physical world? We needn't postulate electrons or sub-atomic particles, but we needn't postulate mice in the walls or birds in the pines, either. If the only reason to accept that something exists is actual, direct observation of the thing, we are reduced to a radical skepticism about the unobserved (but observable) world. As Psillos puts it, "If IBE is generally abandoned, then we are left with a poor epistemology that admits only judgements about observed things. Cartesian skepticism might well be evaded, but Humean skepticism is in the offing" (1997: 371).

van Fraassen embraces this skepticism or, which amounts to the same thing, denies that such a skepticism is a bad thing.[101] His epistemology, recall, consists in the claim that it is rational to maintain unjustified beliefs (as well as justified beliefs), but irrational to not change one's beliefs once they are shown to be unjustified. This matters to the charge of skepticism: the skeptic thinks beliefs about the physical world cannot be justified and so we ought not to hold such beliefs. van Fraassen agrees that such beliefs cannot be justified, but he thinks there is nothing irrational in holding them. Hence his position does not entail skepticism and so the rejection of abduction, both vertical and horizontal, does not result in a *reductio*.

It's not clear that this works, though. Suppose we reject all abduction- not as irrational, but as not rationally compelling. Given van Fraassen's empiricism, it follows that we should reject the conclusions of all ampliative inferences. But this yields the skepticism of which Psillos accuses van Fraassen: we are not rationally compelled to believe in the existence of the physical world. Many would say that this constitutes a *reductio*; belief in

---

[101] van Fraassen (1989) claims that the position does not yields skepticism. Ladyman, et. al., accept that it yields skepticism but deny that this is a bad thing. In the following, I follow van Fraassen (1989: 178). I suspect the difference is merely verbal.

the external world is, they would say, rationally compelling.[102] On this view, van Fraassen can't have it both ways: if we reject all abduction then we must fully accept the consequences of this rejection, and not just the consequences that are consistent with belief in the physical world.

So the rejection of abduction, *simpliciter*, yields global scepticism, which might be viewed as a *reductio* on such a rejection. We needn't enter the debate over whether this is the case, though, because this is the consequence only of the rejection of all abduction. If van Fraassen takes the other option available to him and rejects only vertical abduction (or, on our re-formulation, unfamiliar abduction), he is entitled to his selective skepticism but saved from skepticism about the physical world. He would still have to defend the observable / unobservable distinction, but this he is willing to do. He also would have to defend selective scepticism about the unobservable, but this, too, he is prepared to do. He needn't be a sceptic to undermine realism: he need only be an empiricist.

---

[102] It is worth noting, in this context, that none of van Fraassen's co-authors in the Ladyman article share his belief in the viability of the rejection of all abduction (p. 320).

5. BEYOND REALISM AND ANTI-REALISM

If the analysis of the preceding chapters is right, there is neither a realist nor an anti-realist account of the success of science to be had. At this point, we can either decide, with Popper, that success is inexplicable,[103] or we can look beyond realism and anti-realism for an explanation. What sort of explanation might this be? In this chapter I explore the options for an irrealist[104] account of success- that is, an account that is neither realist nor anti-realist. First, though, I want to review what we have learned about anti-realism; this will allow us to get a better sense of what irrealism is and why it is worth exploring.

5.1 what anti-realism has in common with realism

We noted in chapter one that anti-realism is characterized by the denial of the existence of the entities of science (strong anti-realism) or the denial of the mind-independent existence of those entities (weak anti-realism). This is well and good, but it leaves open the question of what it is to be an anti-realist about science. Does it come to a denial of success? Does it come to a belief in the inexplicability of success? Our chapter one definition circumscribes the view well enough, but it offers a purely negative account of what it is to be an anti-realist. We are left wanting to know what it is, beyond rejecting realism, to be an anti-realist. We have just analyzed two anti-realist attempts at such an explanation; what do they have in common? I want to suggest that, like the realist, the

[103] Cf. Hacking (1983: 57): "Popper ... writes that it never makes sense to ask for the explanation of our success. We can only have the faith to hope that it will continue. If you must have an explanation of the success of science, then say what Aristotle did, that we are rational animals that live in a rational universe."
[104] Nelson Goodman uses this term to describe his radical nominalism, according to which everything – from stereos to stars, from computers to quarks – is made up, the result of the entrenchment of predicates. (See his *Fact, Fiction, and Forecast* and *Of Mind and Other Matters*.) I intend a broader use, indicating

anti-realist believes that the explanation of the success of science is somehow internal to science - that we can look to science to account for its success.

We saw this play out in our studies of van Fraassen and Fine. van Fraassen defends an evolutionary account of success, according to which successful theories survive and unsuccessful theories fail, and this is why our theories are successful. We saw that this account of success is incomplete because it fails to explain why particular theories succeed. Notice, though, that while it is anti-realist – it rejects the explanatory appeal to the ontology of science – the view appeals to the non-ontological resources of theories – their survival – in order to explain their success.

So, too, for Fine. His Natural Ontological Attitude asks us to 'take science at face value', without amplifying our ontological commitments by appeal to a theory of truth or a metaphysics. As van Fraassen suggests an evolutionary approach to success, Fine suggests a naturalistic approach: we should look to the ground-level discipline, and not to philosophical analyses of the discipline, for answers to our questions. Again, there is an appeal to theoretic resources: the ground level commitments of theories explain their success.

Fine and van Fraassen share a negative project – they both refuse to appeal to the ontology of theories to explain their successes – but they share a positive project as well: both think that we can appeal to some non-ontological feature of theories to explain

---

any position that is neither realist nor anti-realist. To whit, I will use the terms 'irrealist' and 'non-realist' interchangeably.

success. Why should we think that this is characteristic of anti-realism generally? I'm suggesting that if you are playing the game with the realist – opposing realism – then you've got to play the game on his terms, which requires keying your account of success to ontology: it's not the ontology of the theory that explains its success, but the _____ of the theory, where the blank is filled in by *some other feature* of the theory. The issue is one of contrastive explanation: the foil of the explanation of success is provided by the realist (the ontology of the theory) and the contrast class is the set of other features of the theory.[105]

If this is right, we can usefully characterize irrealism as a program that seeks to explain the success of science without appealing to features of science. We'll need to see just what this comes to. First, though, let's examine the implications of this analysis of anti-realism.

5.2 a principled objection to anti-realism

I have suggested that van Fraassen and Fine both make failed attempts to explain success in non-ontological, theoretic terms. van Fraassen's attempt fails because neither survival nor empirical adequacy is sufficient to explain why particular theories succeed. The problem with Fine's view is not that it fails to explain success, but that it fails to distinguish itself from realism: the realist says that we ought to be committed to whatever our best theories tell us; the naturalist says that all there is, is what our best theories tell us there is. Save for a nominal difference between 'we ought to be committed to' and 'all

[105] See Lipton (1991:75-98). It's possible to take the contrast class to be something other than features of theories – other ontologies, maybe – but it's not at all obvious how such an account of success would work

there is' – a distinction that makes no difference, to be sure – these are identical commitments.

So NOA collapses into realism because it looks to the ontology of science to explain the success of science. Constructive empiricism does not collapse into realism because it refuses to recognize any feature of science other than empirical adequacy. NOA succeeds in accounting for success; constructive empiricism does not. Notice that what allows NOA this success is precisely what leads to the collapse. Notice also that constructive empiricism cannot obviously be modified to allow it an adequate explanation of success without suffering the same fate, for such a modification would have to identify those features of successful theories that allows them to survive and succeed. But what would those features be? If we appeal to anything like the truth of the statements of the theories or the existence of the posits of the theories, the collapse into realism is immediate. And what other options are there? What features of theories are there that could account for success but do not usher in realism? There simply do not seem to be any: the choices are an unhelpful account of success or a realist account of success. The seeming moral of the story: *anti-realism is in principle an untenable position.*

The point is worth recapitulating. Realism and anti-realism both look to the resources of scientific theories in order to explain the success of those theories. The thrust of argument just rehearsed is that the only theoretic resources that can explain success are ontological resources – this was the thrust of the argument of chapter one – and it is the

---

or how it could be taken to characterize either van Fraassen's or Fine's theories.

realist who looks to ontology. Hence any appeal to theoretic resources must, to account for success, be an appeal to ontology. Thus we see putative anti-realisms collapse into realism, as Fine's does, or resist the collapse but fail to account for success, as van Fraassen's does. If we are to engage the realist, the realist wins.

The point can be put differently, as follows. This is a game of choices. You must choose whether to explain success or not. If you choose to explain success, you must choose whether to explain it in theoretical terms or not; I have suggested that this is the choice between a anti/realism and irrealism. If you choose to explain it in theoretical terms, you must choose whether to explain it in ontological terms or not; I have suggested that this is the choice between realism and anti-realism, and I have further suggested that this is a false dichotomy: there is no non-ontological, theoretical explanation of success to be had. But forget about names of positions; the conclusion we have reached here is that the options close off, and the only viable choices left are explaining success in non-theoretical terms or not explaining it at all.

5.3 sociology of science

The irrealist thinks that we can opt out of the realism game: we can refuse to explain success by appeal to theories. But can we adequately account for the success of science without appealing to theoretical resources? Proponents of the relatively recent tradition of Social Studies of Science think we can.[106] Philosophy of science has traditionally been committed to two methodological tenets. The first, examined above, is internalism: the belief that the explanation of the success of science is to be found within science itself.[107] The second is a commitment to the discovery/justification distinction and, more specifically, to the belief that issues of justification are the only ones that are philosophically interesting. How science yields theories is a matter of science; how theories are justified is a matter of philosophy. The sociology of science seeks to invert this latter commitment and examine the social milieus in which scientific experiments are conducted.

David Bloor, for instance reasoned that, if we consider only the results of science, science appears to be univocal and given to from nature, but if we examine contexts of scientific discovery, we find that those results are won by individuals, sometimes working alone and sometimes in groups, who have their own interests, biases, and limits, whose equipment sometimes malfunctions, who fight with other scientists about how to interpret experiments, etc.[108] Bloor concluded that scientific knowledge is not given to us by nature. By sociologizing science – by treating the work that scientists do as relevant to the theories science produces – Bloor drove a wedge between nature and our knowledge of it: science does not reveal to us the way the world is; rather, it constructs models of the world.

---

[106] There's a tremendous literature: for a nice overview, see Hacking (1999), Kukla (2000) or Galison and Stump (1996). I will use the slightly more manageable 'sociology of science' to name this school of thought.

[107] This connects to the discussion of the previous section: both realism and anti-realism, insofar as they are philosophies of science, look to science itself to explain the success of science.

[108] Bloor developed the Strong Programme in the Sociology of Scientific Knowledge, also known as *the Edinborugh School*, which we will examine in detail a bit later on. I will use *The Strong Programme* to pick out Bloor's theory, and I will use *Sociology of Scientific Knowledge* (SSK) to pick out the several theories which, like Bloor's, examine science from a sociological perspective.

5.4 social constructivism

The sociology of science thus examines the social contexts in which scientific theories are produced. It is out of this methodological tradition that social constructivism is born. Social constructivists take seriously (literally, insofar as it is possible) the conclusion that scientific knowledge is constructed by scientists rather than given to us by nature; theories do not tell us what the world is like. Unlike their anti-realist cousins, however, constructivists go beyond the resources of scientific theories to explain the development and success of those theories: they appeal to the social contexts in which the theories are produced (contexts of discovery).

But what does it mean to say that scientific knowledge is constructed, and how does this explain why theories are successful? Does the social constructivist offer us anything other than the truism that scientific theories are produced in social contexts? Let us look at the main schools of constructivist thought and then, after showing why these programs fail to account for success, move on to an analysis of the problem. This will allow us to identify the desiderata of a successful sociological account of success, and I will then proceed to examine the prospects for such an account.

5.4.1 the Strong Programme

We looked at Bloor's theory a bit in the previous section. It was born of the sociological interest in examining contexts of discovery rather than contexts of justification. This interest underwrites the Strong Programme's *naturalism*, according to which science should be studied no differently than other aspects of human culture, with no preconceived notions of the truth or falsity of any of its claims. In addition to naturalism, the Strong Programme is characterized by a commitment to (what I will call) *externalism*: a causal analysis of scientific knowledge in terms of social conditions. (This stems from the rejection of internalism, noted above as one of the characteristic commitments of the philosophy of science). The Strong Programme is further characterized by a *symmetry postulate*, according to which true scientific theories are to be explained in the same terms as false ones.[109]

The Strong Programme, then, issued a call to examine contexts of discovery as opposed to contexts of justification (naturalism), and to pay attention to the social aspects of those contexts (externalism). It also called for symmetry in accounting for the development of science: previously, sociological analyses had been afforded a role only in the analysis of false theories: when science worked, it was because reality made it so; when science failed, it was for sociological reasons. Naturalism rules out this *a priori* sorting of explanations and it, combined with externalism, leads to the rejection of realism's "*a priori* assumption that it is reality that fixes beliefs" (Nanda 7).

While he rejected this assumption, however, Bloor was unable to give a clear account of how social factors are involved. In addition, he allowed that, in fact, social reasons play only a partial role in fixing belief, thereby making room for the realist to re-assert the primacy of reality. That is, Bloor argued that social reasons are involved in both successful and failed theories – he adhered to the symmetry postulate – but he afforded

[109] It is also committed to a reflexivity postulate, according to which the analysis of science offered by the sociology of science should also apply to the sociology of science itself. This doesn't come into play in the

them only a limited role. In so doing, he made room for a host of other, stronger versions of his Strong Programme.

5.4.2 the Bath School

Proponents of the Bath School – mainly, Collins and Pinch – also advocate the principles of symmetry, naturalism, and externalism. Hence they, too, reject realism's *a priori* assumption that reality fixes scientific belief. Where Bloor would have us look to the social aspects of the context of discovery, generally, however, the Bath School would have us look at contexts of discovery on the micro-level: on the level of individual activities of individual scientists engaged in scientific disputes. This is a part of what is known as 'controversy studies': a sociological program of analysis of scientific controversies designed to identify sociological contributions to the progress and development of scientific knowledge. Specifically, Collins and Pinch argue that scientific controversies are closed (resolved) by a host of mechanisms external to the contents of the controversies: "faith in experimental capabilities and honesty..., personality and intelligence of experimenters, reputation of a running lab, whether the scientist worked in industry or academia, previous history of failures, 'inside information', style and presentation of results, psychological approach to the experiment, size and prestige of the university of origin, integration into various scientific networks, nationality" (Collins 1992: 87).

This is an extension and refinement of Bloor's Strong Programme: it pushes naturalism and externalism into the laboratory, as it were, and forces the sociological examination of

---

present analysis.

contexts of discovery to its natural conclusion. Because so many of the factors involved in resolving disputes are fluctuant, however, it follows that the resolutions of controversies themselves can fluctuate. Less obtusely, the resolution of any given scientific controversy is a highly contingent affair: were we to change the nationality, say, of a key player, or her university affiliation, a given controversy might have a radically different resolution—that is, were we to alter these sociological variables, we might well alter the results of scientific experiments. Unlike Bloor, then, Collins and Pinch invite a far more radical irrealism about science.

5.4.3 Discourse Analysis

Discourse Analysis takes micro-history and controversy studies to the next level: rather than examining the disputes of scientists, it examines the discourse of the scientists engaged in those disputes, where 'discourse' includes to not just verbal discourse, but also written discourse ('inscriptions'). The governing idea here is that the currency (as it were) of science is the written report –lab reports, monographs, journal articles, etc. – and so the factors that enter into the analysis of these written reports is central to the analysis of science. Discourse Analysts (if we may call members of this school that) argue that, much as they do for Bloor and Collins and Pinch, external factors contribute to the creation and assessment of these written reports and, hence, that external factors affect in a serious way the results of science.

This is a further extension and refinement of the Strong Programme. Picking up on Bloor's failure to explain how social factors are involved in science, Discourse Analysists offer an account of how this works. Scientific results –reports, articles, books, etc. – are finished products, the earlier drafts of which we never see. We grow accustomed to taking the results of science at face value – to never looking at the work that went into shaping them – and so to accepting the results of science on faith. If we look at that earlier work, however, we see that a host of considerations external to the content of science affect the production and acceptance of scientific papers: the need to cite authorities, the desire to have one's own paper cited, etc. (Golinski 37-44).[110] Since these papers present the results of science, it follows that external factors affect the results of science. Hence Bloor's initial sociological program is realized in the micro-history of science.

5.4.4 failures of these programs

It should be fairly obvious that none of these programs stands a chance of accounting for the success of science. All three offer analyses of the production of scientific knowledge – analyses of how scientists behave and what influences their decisions – but none of them explains why theories are able to yield successful predictions. These programs offer sociological accounts of theory acceptance, but not of success. Bloor's Strong Programme, for instance, allows that true scientific beliefs (that is, belief in true scientific theories) are to be explained in the same way as false beliefs, and that the explanans in both cases should be sociological. But this does not explain why some theories are true (succeed) and others are false (fail); it offers only a sociological account of why scientists choose those theories and not others. Similarly, the Bath School's call to examine

[110] Latour (1986) grants that experimental results are also important, but he argue that experiments stand in need of interpretation and analysis and so experimental results can themselves be subjected to a discourse

scientific controversies and to acknowledge the role of the social in the resolution of those controversies fails to identify any features of scientific theories in virtue of which those theories issue in successful predictions and explanations. Discourse Analysis fails for precisely the same reason. In the end, these are all studies of how scientists behave and, as Kukla (2000: 43) notes, "it shouldn't come as a surprise that questions about the nature of the world and the status of science can't be settled just by looking at how people behave."

5.5 constructivism about what?

The failures of the main forms of constructivism lead us to ask what constructivism must do in order to account for the success of science and so be a viable alternative to realism (and anti-realism). This leads us to examine the several possible objects of constructivist analysis. Hacking (1999: 21-22) notes that one can be a constructivist about objects (X), ideas (the idea of X), or what he calls elevator words, or semantic claims about X (facts, truth, reality and knowledge).[111] Let's begin with knowledge.

5.5.1 epistemic constructivism

Social constructivism grew out of the examination of contexts of discovery in order to conduct a sociological investigation of the genesis and development of scientific knowledge. This led to the study of scientists and to the belief that scientific knowledge is constructed (read: not fixed by reality). Call this *epistemic constructivism*; this is the

---

analysis.

[111] It's odd that Hacking includes *knowledge* in his list of elevator words because, while *reality*, *fact*, and *truth*, all seem semantico-metaphysical, *knowledge* is semantico-epistemological. Knowledge claims might be contingent in all the ways the constructivist says while the thing itself (and perhaps the idea of the thing)

heading under which the major theories examined above can be placed. Jan Golinski characterizes epistemic constructivism as the view "that scientific *knowledge* is a human creation, made with available material and cultural resources, rather than simply the revelation of a natural order that is pre-given and independent of human action" (1998: 6, my emphasis). He makes clear, however, that this is not to be confused with idealism: constructivism "should *not* be taken to imply that science can be entirely reduced to the social or linguistic level, still less that it is a kind of collective delusion with no relation to material reality" (6). For Golinski, constructivism "is more like a methodological orientation than a set of philosophical principles; it directs attention systematically to the role of human beings, as social actors, in the making of scientific *knowledge*" (6, my emphasis).

Golinski's account is explicitly keyed to knowledge; the claim is that *scientific knowledge is a human construction*. As we have seen, this is too weak a conception of constructivism to underwrite an account of success. A parallel to our analysis of the several kinds of realism is instructive: the epistemic conception of realism faltered because the fact that our belief in a theory is justified fails to explain the success of that theory. Similarly, the fact that our belief in a given theory is the result of social agency (however that works) cannot explain why that theory issues in successful predictions.

---

are not. We'll see in a few paragraphs that all this comes to a quick end for the prospects of an epistemic constructivism.

5.5.2 semantic constructivism

5.5.2a truth

Hacking's list of elevator words includes knowledge, truth, facts, and reality. We have seen that constructivism about knowledge cannot explain the success of science. What about truth? To be a constructivist about truth is to claim that the truth of some sentence X is not independent of humans but is, rather, somehow created by us. Of course, since sentences are human constructions, we expect humans to be involved in the production of true sentences. Our analysis here mirrors the analysis of the semantic conception of realism: a theory about words cannot account for the predictive success of scientific theories. This construal of constructivism amounts to the following: claims about scientific facts of the matter are not inevitable; rather, they are shaped by social events and forces and so might well have been different. No surprise here; not even the most devout realist suggests that there is an inevitability to what we say about the world.

But one might think that the truth conditions for sentences are independent of humans. I take the constructivist to be denying this claim: we create the conditions under which a given sentence is true. Here, too, our analysis mirrors the analysis of the semantic conception of realism. If the constructivist is claiming that the truth conditions for a given sentence are the result of human activity, he is making a claim about the world, not about words. He is saying that our sentences are true in virtue of the fact that they accurately represent a reality that is (in some way) constructed. That's well and good, but that's a constructivism about reality, not truth.

5.5.2b reality

Reality is on Hacking's list. What does it mean to say that reality is constructed? Hacking has in mind the epithet 'the reality of x', as in 'the reality of serial killing', or 'the reality of neutron stars' (22-24). Clearly, this reduces to ontology: to say that the reality of X is the result of human activity is simply to say that X is the result of human activity. Like 'the truth of', 'the reality of' is a straightforward case of semantic ascent and, as such, lends no non-ontological hand to the explanation of success.

5.5.2c facts

What about facts? *Prima facie*, it sounds reasonable to account for the success of a given theory by appealing to the claim that the fact that X is constructed. (Indeed, Latour and Woolgar's seminal *Laboratory Life* is subtitled '*The Social Construction of Scientific Facts*.') But what does it mean to say that a fact is constructed? The question is: is there any difference between 'the fact that X is constructed' and 'X is constructed'? Are facts, too, parasitic on ontology? In one sense, yes: if the fact in question is the fact of the existence of X, then there is no difference. But suppose we are talking about an event – a vapor trail, say, left in a bubble chamber. We might say that the vapor trail is constructed, and this would reduce to ontology. But we might also say that the vapor trail was left by a passing electron, in which case the fact in question would be the electron leaving the vapor trail. Here we're not just talking about objects; we're talking also about a causal relationship, and about the event that took place- the fact that the electron left a vapor trail. Constructivism about the causal relationship, like constructivism about the vapor trail, reduces to ontology (though process, not entity). But the event – the leaving of the trail –is neither entity nor process.

So here we have a kind of semantic constructivism that does not seem to reduce to ontology. But can it explain success? The constructivist about facts is saying that we can characterize goings-on in the world in different ways, and these characterizations are the result of social human activity. No surprise, when it's put like that. But, the realist is sure to respond, the fact that a theory characterizes an event in one way rather than another cannot explain the theory's success. That Michelson and Morley got a null result has nothing to do with their characterizations of their experiment; it is explained by the fact that there is no aether.

We've seen this knee-jerk realism before, though. The assertion of the realist explanation of success is intuitively appealing, but flawed. Surely, the constructivist can respond as follows: how you characterize a situation has got lots to do with whether your theories about the situation succeed. If you characterize the vapor trail as moving up rather than down, your theory had better account for that if it is to be successful. If you characterize a detector event as neutron event rather than, say, as an electron event, a neutron theory will succeed and an electron theory will fail. To be sure, our characterizations of events will be constrained by the way things are, but this does not show that characterizations do not matter.

Constructivism about facts, then, unlike constructivism about truth, knowledge and reality, does not simply make a semantic ascent and so does not immediately reduce to ontology. So here we have a desideratum on constructivism: it cannot take as its object knowledge, reality, or truth, and it can take facts to be constructed.[112]

5.5.3 constructivism about ideas

What does it mean to talk about the idea of X being the result of human activity? Hacking has in mind "ideas, conceptions, concepts, beliefs, attitudes to, theories" (22). The constructivist about ideas, then, might be saying that our concepts and beliefs are the result of human activity. But this can't be right; even the most devout realist will admit that, without humans, there are no beliefs. This is neither enough to account for the success of science nor enough to underwrite a substantive alternative to realism. Nor could the constructivist reasonably be taken to be saying that the causes of our beliefs and ideas are constructed; this would straightaway reduce to constructivism about objects. Moreover, how could the idea of a quark – when explicitly divorced from any ontological cause of the idea – possibly explain the success of quark theories?

So constructivism about ideas cannot be merely epistemic. Notice that Hacking takes 'idea' to include concepts and theories. Perhaps, then, constructivism about ideas is more akin to constructivism about facts: perhaps the constructivist about the idea of X is saying that the categorization of the thing as an X-thing is not given to us by nature but is, rather, the result of human activity. This conglomeration of people is a theater audience, for

[112] That is, we can make sense, in way that we can't for elevator words, of talk of construction of facts explaining success. But this does not pose any special difficulties for the analysis of constructivist accounts

instance, because it satisfies certain, artificial conditions; if different conditions obtained, it would be a mob or a panel. This is a familiar line of argument: what makes something a chair is the use to which it is put.[113] It is a small step from this to a constructivism about social objects – what makes something an audience is the set of criteria satisfies – and then to physical objects: what makes this arrangement of particles a particular (scientific) thing is the fact that it satisfies certain criteria.

Can this explain success? It seems a likely candidate—but only because it tells us how things are constructed. That is, if it is able to explain the success of, say, the theory of quarks, it is because it tells us how quarks – the things themselves – are the result of the imposition of ideas on the world. It's not the idea that is constructed; it's the thing. Constructivism about ideas tells us how constructivism works – it gives an account of concept imposition – but it doesn't tell us what is constructed.

5.5.4 ontological constructivism

We keep coming back to constructivism about things. This makes sense; the upshot of our chapter one analysis of realism was that non-ontological versions of realism were either parasitic on the ontological version or were unable to account for success. So, too, it seems, for constructivism.

But what does it mean to say that things are constructed? Recall Devitt's analysis of realism in terms of the dimensions of existence and independence: you can reject realism

---

of success; at most, it will require us to complicate our ontological categories a bit.

[113] See Wittgenstein, of course, and Margaret Gilbert.

by either denying that a thing exists or by denying that it exists independently of humans. The constructivist denies that the constructed entity has an independent existence; it's existence is somehow contingent on the existence of humans. We can make sense of this for artifacts- things that are literally constructed: the bookcase that I assemble in my garage is constructed; without me, it would not exist. To think that, say, quarks are constructed is to think the same thing about quarks: without humans, they would not exist. Of course, I do not put together quarks in my garage. But, the constructivist maintains, scientists put them together (as it were) in their laboratories. The difference is that quarks are not, *prima facie*, artifacts. Rather, we take them to be natural kinds- to exist independently of humans. And this is what ontological constructivism, in the end, comes to: the constructivist maintains that things commonly thought of as not being the result of human activity really are those sorts of things.

Suppose we can make sense of this. Can constructed entities and processes account for success in the same way as natural entities and processes? The realist said that the reason there happens what the theory says will happen is that the world is the way the theory says it is; the constructivist says that there happens what the theory says will happen because there is what the theory says there is. The difference is that, for the realist, reality precedes theory and, for the constructivist, theory precedes reality: the constructivist thinks that we create (in a way to be described) the entities and processes postulated by theories and they explain success, while the realist thinks that theories describe an independently existing reality. A substantive difference, to be sure, but the explanatory point remains: if sense can be made of ontological constructivism, it can account for success.

5.6 what does it mean to say things are constructed?

Our analysis thus far mirrors our analysis of realism: epistemic and semantic versions are either parasitic on ontological versions or are unable to account for success. What we need to understand now is what it means to say a thing is constructed.

On Ian Hacking's analysis, to be a constructivist about X is to hold that "(1) X need not have existed, or need not be at all like it is. X, or X as it is at present, is not determined by the nature of things; it is not inevitable."[114] He notes that constructivists also often hold that "(2) X is quite bad as it is" and that "(3) We would be better off if X were done away with, or at least radically transformed" (6). Clearly, (2) and (3) are extra-metaphysical; they are remnants of the political uses to which constructivism is often put.[115] A constructivist about quarks, for instance, is not likely to hold theses (2) and (3); indeed, it isn't clear what it would even mean for to think such things about quarks. The central thesis of constructivism is what Hacking calls the inevitability thesis: "The existence or character of X is not determined by the nature of things. X is not inevitable. X was brought into existence or shaped by social events, forces, history, all of which could well have been different." (6-7).

[114] All references to Hacking are to Hacking (1999) unless otherwise noted.
[115] Hacking acknowledges this; see pp. 7-9.

So the constructivist thinks that things thought to have an independent existence really do not. We still don't know what it means to account for something's existence in social terms. We have seen, however, that it makes sense to think of it in terms of the imposition of categories on the world. Michael Devitt (1991: 235) describes constructivism as a combination of Kantianism and relativism, characterizing it as the view that "the only independent reality is beyond the reach of our knowledge and language. A known world is partly constructed by the imposition of concepts. These concepts differ from (linguistic, social, scientific, etc.) group to group, and hence the worlds of groups differ. Each such world exists only relative to an imposition of concepts." Devitt is here after Nelson Goodman's radical irrealism, which takes as its object potentially anything and virtually everything. That's more than the contextualist need be committed to,[116] but Devitt is onto something. We have characterized constructivism about X as the view that X is the result of (shaped by, tempered by, affected by) human activity and is not, as is thought, mind-independent, and we have noted that an account of human activity is needed. Devitt suggests that the activity in question is the imposition of concepts. This sounds right: if we trace the history of constructivism back through history, we cut an arc through Kuhn to Kant and Berkeley- and each of these theorists suggested that the imposition of concepts on the world accounts for our understanding of the world. The particulars of Goodman's view – that concept imposition is a socio-linguistic activity in that concepts are imposed by individuals in virtue of their belonging to particular linguistic groups – can be left aside.

---

[116] One is reminded of Michael Levin's (1990) argument that one has to be a realist *about something*. So, too, for constructivism.

So the social constructivist maintains that scientific entities and facts are not given to us by nature but are, instead, brought into existence by the imposition of concepts. If this is the case, then the desiderata on an account of constructivism are clear: it must offer an account of concept imposition (including an account of how concept imposition yields scientific facts/things) and it must show how the imposition of concepts explains success.

5.7 Pickering's *opportunism in context*

Can these desiderata be met? Andrew Pickering's *opportunism in context* points the way towards a version of constructivism that might do the trick. Pickering contrasts the scientist's account of the history of the quark with the historian's.[117] According to the former, "experiment is seen as the ultimate arbiter of theory. Experimental facts dictate which theories are to be accepted and which rejected" (Pickering 1984: 5). The historian's account stems from "two well-known and forceful philosophical objections to this view, each of which implies that experiment cannot *oblige* scientists to make a particular choice of theories" (5). The first objection is the thesis of underdetermination. The second objection – and this is the one that leads to Pickering's constructivism – "is that the idea that experiment produces unequivocal fact is deeply problematic" (6). It is problematic, Pickering thinks, because experiments are performed on "'open', imperfectly understood systems, and therefore experimental reports are *fallible*" (6). They are fallible because reports of experimental outcomes depend on the theories that

[117] Pickering (1984) is concerned specifically with the history of High Energy Physics, and so he casts his account in terms of that history. His account serves well as an account of the history of science, generally, though, and the reader should regard it accordingly.

All references to Pickering are to Pickering (1984) unless otherwise specified.

inform the experiments, and if these theories change then the experimental outcomes will change. Moreover, experimental reports are fallible because, although

> Experimenters do their best ... to eliminate all possible sources of background, ... it is a commonplace of experimental science that this process has to stop somewhere if results are ever to be presented. Again, a *judgement* is required, that *enough* has been done by the experimenters to make it probable that background effects cannot explain the reported signal, and such judgements can always, in principle, be called into question (6).[118]

Thus we arrive at what Pickering calls the historian's view (and what we call constructivism): he argues that the social contexts in which these judgements are made play a significant role in determining experimental outcomes. Specifically, Pickering maintains that "one can only appeal to the reality of theoretical constructs to legitimate scientific judgements when one has already decided *which* constructs are real. And consensus over the reality of particular constructs is the outcome of a historical process" (7).

Clearly, Pickering is interested in ontology and not mere knowledge. His claim is that we determine what exists on the basis of experimental results, but these results are themselves the result of an historical process- the judgements of experimenters that their theories are adequate and that they have done enough to eliminate the background. To be

[118] This is a variant on the experimenter's regress (Collins 1992). Given an experimental set-up and a set of experimental results, Collins suggests, the experimenter must determine what counts as a good result and what counts as a competent experimental set-up. But this is not easily done: knowing what counts as a good result presupposes knowledge of what counts as a competent set-up, and knowing what counts as a competent set-up presupposes knowledge of what counts as a good result. If an experimenter knows what the results of the experiment should be, he could use this knowledge to determine when the set-up is flawed; if the experimenter knows which experiments are competent, he could exclude results accordingly. Collins claims that scientists know neither authoritatively, and he suggests that (as a matter of fact) they resolve the dilemma by determining the competency of experiments on the basis of their expectations of the outcomes.

sure, Pickering writes: "the scientist legitimates scientific judgements by reference to the state of nature; I attempt to understand them by reference to the cultural context in which they are made. I put scientific practice ... at the centre of my account, rather than the putative but inaccessible reality of theoretical constructs" (8). This talk of the putative but inaccessible reality of theoretical constructs, combined with the claim that we appeal to the reality of theoretical constructs to legitimate scientific judgements after we have decided which constructs are real, paints a thoroughly non-realist picture: scientists decide what exists and then attribute reality to those posits. Alternatively, the entities of science are theoretical constructs.

That's Pickering's view, anyway. How is it supposed to work? What we get from Pickering is a short account of the theory and a long case study. The theory is called *opportunism in context* (OIC). It "seek[s] to explain the dynamics of [scientific] practice in terms of the contexts within which researchers find themselves, and the resources which they have available for the exploitation of those contexts" (11). The 'context' to which Pickering refers has a justificatory element and a contentful element. These elements are mutually reinforcing:

> to justify his choice to work within [a research tradition], a theorist has only to cite the existence of a body of work of experimental data in need of explanation. And fresh data, from succeeding generations of experiment, constitute the subject matter for further elaborations of the theory. Conversely, for the experimenter, his decision to investigate the phenomenon in question rather than some other process is justified by its theoretical interest, as manifested by the existence of the theoretical tradition. ... Thus, through their reference to the same natural phenomenon, theoretical and experimental traditions constitute mutually reinforcing contexts (10).

The problem of fallibility – the essential role of judgement in experimental results – serves to eradicate the scientist's picture of science as evidence-driven; when experiment fails to "conform to prior expectations … [experimenters] are then faced with one of the problems of scientific judgement noted above, that of the potential fallibility of all experiments. Have they discovered something new about the world or is something amiss with their performance or interpretation of the experiment? From an examination of the details of the experiment alone, it is impossible to answer this question" (9). The symbiosis of justification and content – OIC – then comes into play:

> Now suppose that a theorist enters the scene. He declares that the experimenters' findings are not at all unexpected to him – they are the manifestation of some novel phenomenon which has a central position in his latest theory. This creates a new set of options for research practice. First, by identifying the unexpected findings with an attribute of nature rather than with the possible inadequacy of a particular experiment, it points the way forward for further experimental investigation. And secondly, since the new phenomenon is conceptualised within a theoretical framework, the field is open for theorists to elaborate further the original proposal (9).

On Pickering's idealized picture, the failure of the original experiment, when viewed in light of the new theory, is not a failure at all; rather, it is a confirming instance of the new theory. Further experimentation proceeds and the new theory is vindicated, ontology and all. In short, theories are not developed in response to experimental facts of the matter but, rather, are fitted to experimental situations. It follows that our ontologies are not 'given to us' by experiment; they are the result of this 'fitting' of theory to fact. Hence we get an ontological constructivism that promises to meet our ontological desideratum and so form the basis of a constructivism account of success.

I say that OIC promises to meet our desideratum (and not that it does meet it) because, while it provides an account of concept imposition, it is an account of theory choice, not success. OIC shows us that some version of constructivism might be viable, but it does not show us that constructivism can account for success.

6. BEYOND REALISM AND ANTI-REALISM, PART II: THE PROSPECTS FOR A SOCIOLOGICAL ACCOUNT OF SUCCESS

We have seen that the only version of constructivism that stands a chance of accounting for success is a constructivism about ontology, and we have seen that Pickering's OIC is such a constructivism. It does not follow, of course, that OIC does account for success: satisfying our ontological desideratum is a necessary but insufficient condition on a constructivist account of success. What else is needed? An explanatory mechanism: a theoretical description of how it is that theories are able to accurately tell us what is going to happen. In the remainder of this essay, I explore the prospects for such a description. After sketching what seems to me to be the most likely one, I offer reasons to think this account fails to do the job.[119]

6.1 opportunism in context, again

Pickering's theory, as described, offers only an account of explanatory success- of how it is that theories are able to accurately tell us what has already happened. As we have seen, this isn't too impressive: given that explanations can be fitted to the facts, explanatory successes are relatively easy to come by.[120] But Pickering's theory seems to contain the germ of a more ambitious account of success. Pickering argues that experiments are fallible, that this fallibility is due in part to the role of judgement in determining when enough has been done to rule out false causes, and that there is a symbiosis of content and theory that is in play in making this determination.

---

[119] Full disclosure: it took much beating about the head and shoulders for me to come to this conclusion. John Greenwood and Michael Levin can take credit for most of these beatings: Arnold Koslow can take credit for a kinder, gentler persuasion. Thanks to all.

[120] See chapter two. Indeed, the reason we offered for this is awfully similar to what Pickering maintains is the case: theories can be fitted to the facts (after the fact), and so explanatory success is cheap.

Throughout this analysis, Pickering is concerned with the *ends* of experiments. But what if we were to iterate his analysis throughout the experimental process? This would amount to the claim that, at *any* stage of experimentation, when experiment fails to "conform to prior expectations … [experimenters are faced with the problem] of the potential fallibility of all experiments. Have they discovered something new about the world or is something amiss with their performance or interpretation of the experiment? From an examination of the details of the experiment alone, it is impossible to answer this question" (9, previously cited). In short, we might claim that the mutual reinforcement of theory and evidence is present at each stage of the experimental process (and not just at the end stage of a failed experiment). On this view, all experimental results are viewed by the lights of particular theories, and these theories determine what counts as an acceptable result. If the evidence substantiates the theory, the experiment is viewed as a success. If the evidence does not substantiate the theory, the theory, the apparatus, and the (interpretation of the) results are tweaked until they come into resolution or until the theory otherwise ends.[121]

6.2 contextualism

Call this extension of Pickering's program *contextualism*. Contextualism is a candidate for a constructivist account of success. According to it, a scientist's theoretical dispositions incline him to think that, say, neutrinos exist (or don't exist), and this ontological disposition leads him to interpret the evidence in particular ways.

---

[121] This is intentionally vague. One problem with the sociological account of success is the problem of accounting for the occasional failures of theories. I consider this problem later (see section 6.4). For now, I want only to make clear that the sociologist (on my interpretation, anyway) does not assume that experiments always succeed- hence the vague 'otherwise ends'.

Specifically, it affects the determination of which experimental effects and results are caused by systematic error, and which are caused by the operation of the sought-after entities. In this way, the experimental process is theory-infected and, since accounts of experimental results are based on these attributions of causes and errors, experimental results are theory infected. This accounts for why theories manage to correctly predict and explain: the explanans and the explanadum (and the apparatus) are adjusted until they come into line.[122]

This sounds like what, in chapter two, we called 'fudging' the data by glossing the facts to fit the theory. There's a difference, though: when the theorist fudges, he conducts an experiment and then consciously adjusts the results to get a desired outcome. The present suggestion is that such adjustments occur before and during the process of experimentation, and often unconsciously.[123] The claim is not that scientists massage data to get particular outcomes; rather, their interpretations of the evidence – which yield the data – are affected by their theoretical commitments: the scientist "fights the systematics until he or she gets the 'right' answer (read 'agrees with previous experiments')."[124] The process to which the contextualist is adverting occurs before the data are generated.

---

[122] Or otherwise end; see previous note. I shall henceforth ignore this qualification.

[123] Michael Levin is concerned that this (and related) "talk of psychological set … ignores the expectation-independent logical implications of theories." I will address this concern in a moment. But the sort of response the contextualist will offer is fairly obvious: while the contextualist allows that theories have logical implications, the interpretation of results – the determination of whether a physical event confirms some or other implication – is subject to interpretation. More on this later.

[124] Physical Data Group, "Review of Particle Properties," *Review of Modern Physics* 52 (1980), p. S286. Quoted in Franklin (1986: 236).

In short, the contextualist thinks that the realist's "picture of a quasi-logical deduction of a prediction, followed by a straightforward observational test is simply wrong" (Collins and Pinch 1993: 52). Rather there is mutual reinforcement of theory by data and of data by theory: as scientists come to believe a theory, they come to see the evidence as supporting the theory; as they come to see the evidence as supporting the theory, they come to believe the theory. Writing about Einstein's explanation of the 'red-shift', for instance, Earman and Glymour (1980: 85) note that

> There had always been a few spectral lines that could be regarded as shifted as much as Einstein required; all that was necessary to establish the red-shift prediction was a willingness to throw out most of the evidence and the ingenuity to contrive arguments that would justify doing so. The [Eddington] eclipse results gave solar spectroscopists the will. Before 1919 no one claimed to have obtained spectral shifts of the required size, but within a year of the announcement of the eclipse results several researchers reported finding the Einstein effect. The red-shift was confirmed because reputable people agreed to throw out a good part of the observations. They did so in part because they believed the theory; and they believed the theory, again at least in part, because they believed the British eclipse expedition had confirmed it.[125]

The point is that evidence does not simply support or refute theory. Rather, evidence has to be interpreted, and interpretation is guided by theory: interpretation by Einsteinian lights allowed scientists to say that Einstein predicted the red-shift, but interpretation by Newtonian lights did not. The realist will be quick to insist that this is a matter of theory acceptance, not success, and that there is a fact of the matter about whether Einstein or Newton was right. The sociologist will disagree: the decision to interpret the evidence by the lights of one theory rather than another theory is itself theoretically conditioned. For the contextualist's part, he will demur from such broad discussions of theory choice and argue only that interpretation of evidence is theoretically conditioned and, as a result, theoretic success and failure are themselves theoretically conditioned.[126]

A comparison to Kuhn might help clarify matters. According to Kuhn, scientists 'see the facts' in one way rather than another, just as viewers who see the picture as a picture of a duck see the lines one way, and those who see it as a rabbit see the lines another way. As a result, differing theoretic pre-dispositions can incline scientists to see the same evidence in different ways. The account sketched here attributes to theoretical pre-dispositions the same power to affect interpretations of experiments, but it locates the effect of that power differently. Kuhn contends that theory infects experiment at a foundational, pre-experimental level: we cannot help but see what theory tells us exists; given a particular paradigm, certain entities and processes simply are a part of our ontology. Contextualism holds that theory disposes the scientist to think there exists the entity in question and to interpret the data so as to substantiate that belief. For the Kuhnian, the influence of theory occurs before experiments begin; for the contextualist, it occurs while experiments are underway.[127]

---

[125] Quoted in (Collins and Pinch 1993: 53). Collins and Pinch draw from this analysis a conclusion similar to the one the contextualist draws from it. They note, however, that Earman and Glymour's own conclusions are far more inclined to realism.

[126] It's worth noting, I think, that this is in part what makes sociological analyses non-realist rather than anti-realist. It might well be the case that Einstein got the world right and Newton got it wrong; for the sociologist, this question is distinct from the question of why scientists choose the theories they do and get the results they do.

[127] And fudging takes place after experiments end. One thing that sets contextualism apart is the fact that it takes actual experimentation as its locus of analysis, not experimental reports (as the fudging experimenter does) or worldviews (as the Kuhnian does).

That said, it is important to note that contextualism is not a species of idealism.[128] It is not the mere psychological set of the experimenter that determines the outcome of an experiment. Nor is it the beliefs of the scientist that determine whether neutrinos (e.g.) exist. Rather, the contextualist maintains that scientific method is more complicated than the realist allows- that theory and data affect each other and that, in order to account for the success of science, we have to appeal to the activities of scientists as well as to an objective reality.[129] This is important: if one loses sight of this basic methodological premise, the theory will never get off the ground- if the reader insists on granting (theoretic) explanatory efficacy only to the independent world and never to human agents, the view simply will not make sense. Again, one need not endorse idealism in order to countenance the view described here: the contextualist does not resist the claim that there is a world independent of humans. What he resists is the claim that this world determines (on its own) the content of our theories. Sociology of science is not a metaphysics; rather, it is a meta-science: it purports to tell us how our theories work, not what the world is like.[130] On the sociological theory developed here, our theories work because theoretic assumptions inform experiments and so affect attributions of cause and error in such a way that theoretic expectations and experimental results often come into line with one another.

---

[128] *Pace* Michael Levin.

[129] Of course the activities of scientists matter in the trivial sense that, if scientists didn't conduct experiments, it would make no sense to talk about the results of experiments. The claim here is intended to be stronger than this; the contextualist claims that the decisions scientists make and the activities in which scientists engage affect the outcomes of experiments. The connection between theory and world is thus mediated by scientists' beliefs and choices.

[130] Arnold Koslow questions this distinction. The realist is likely to reject it, claiming that our theories work precisely because they tell us what the world is like. The sociologist, as is made clear in the text, endorses a substantive distinction between the way the world is and the way theories work.

6.2.1 an example: gyromagnetism

In 1915 Einstein and de Haas conducted experiments to determine the gyromagnetic ratio (the ratio of angular momentum to magnetic moment). They did this for two reasons. First, they had a theoretic interest in confirming Ampere's hypothesis that magnetism is the result of electricity in motion. Second, they had a practical interest in engineering: Einstein was working on the design of a reliable gyroscope, which led him to consider micro-scale versions of the gyromagnetic effect. The experimental goal was simple: they wanted to magnetize a suspended iron rod and thereby set it in motion. Informing their work was a belief that the electron, introduced ten years earlier, was the medium through which Ampere's electromagnetism worked: they wanted to show not only that Ampere's hypothesis was correct, but that the effect was due to electrons.

Einstein and de Haas' belief that bound, orbiting electrons were the source of magnetism led to certain experimental biases in their work. Specifically, it led them to assume that what they were examining was charge-to-mass ratio, and this led them to set certain parameters on the possible results of the investigation. As Galison notes, "Einstein's concerns about the zero-point energy and atomic structure ... led [him] to believe that *g* [the gyromagnetic ratio] ought to be not just of order unity, but 1.0, and not 2, 4, or 1/2. The difficulty then lay in extracting a comparably precise result from experiment" (69). Such extraction consisted in the acceptance and dismissal of evidence; a successful experiment was one in which the data fit with the expected result, while a failed experiment was one in which the data did not fit. To whit, Einstein and de Haas conducted no fewer than four different versions of the experiment, stopping only after they had "convinced themselves that they had confirmed Ampere's hypothesis..." (Galison 47).

But does this show that the success of the theory was due to the scientists' expectations? One can fairly ask: 'Einstein and de Haas's predispositions notwithstanding, their experimental values still coincided with their theoretical values. How is this possible?' The contextualist claim is that the coincidence of experimental and theoretical values occurred because the theoretical value was known in advance and so determined the experimental value. As Galison notes, "Anyone trying the gyromagnetic experiment under the supposition of orbiting electrons would dismiss an effect if it were so large as to be consistent with direct coupling of the magnetized cylinder to the earth's magnetic field" (73). Einstein and de Haas knew what they wanted the measurement of *g* to be, and they ran multiple experiments before succeeding. The stopping point was, naturally enough, the point of coincidence. Einstein and de Haas had in mind what would count as a successful experiment: it would be one in which the measure of *g* fell within their margin of error. As each experiment produced a measurement that was not within the margin of error, they made adjustments and tried again. They stopped when they succeeded.[131]

[131] Writing about a different example (Eddington's eclipse experiments), Collins and Pinch make a similar observation: "Even to have the results bear upon the question it had to be established that there were only three horses in the race: no deflection, the Newtonian deflection, or the Einsteinian deflection. If other possible displacements had been present in the 'hypothesis space' then the evidence would be likely to give stronger confirmation to one or the other of them" (1993: 50). Their point, like mine, is that experimental results are ambiguous, and the bearing of experimental results on theories is equally ambiguous. In order to determine the relevance of experimental outcomes for theories, the evidence must be interpreted, and these interpretations are affected by the decision to consider some hypotheses but not others. The present point is that this ambiguity and interpretation begin when the experiment begins.

So what explains the coincidence of theoretical value and experimental value? The realist contends that the theory 'gets right' a fact about the world, and it is this 'getting right' that explains how the theory is able to account for the results of the experiment. On the contextualist view, the coincidence of theory and experiment is due to the fact that the theoretical value was sought after in the experiments, and the experiments were conducted multiple times, each time being altered in an attempt to bring them more in line with the desired result. The realist contends that the coincidence of theoretical value and experimental value – the success of the theory – is due to the coincidence of postulated ontology with actual ontology; the contextualist account of this coincidence requires no such ontological coincidence, and so it serves as a non-realist alternative account of success.

6.3 the problem of quantitative vs. qualitative success[132]

It would be nice if contextualism were able to account for success. Beyond providing a solution to the problem of success, it would allow us to account for the fact that false theories often seem to be explanatorily successful. As we saw in previous chapters, this is a bit of a sticky wicket for the realist: he must argue either that theories with non-referring central terms nonetheless refer or that seemingly successful theories are not really successful at all. I argued earlier that both moves are convoluted.[133] On the

[132] Thanks to John Greenwood for these categories, and thanks to Greenwood and to Michael Levin for pushing me to acknowledge this problem for contextualism.

[133] I have called the former strategy the 'revisionist reference' strategy and argued that it causes more philosophic problems than it solves. We can call the latter strategy the 'retroactive realist' strategy for, according to it, we cannot determine whether a theory actually explains until we can determine whether that theory is true. Since this can take hundreds of years (if it can be done at all), the strategy requires a sort of eternal wink-and-nod explanatory stance: 'X explains Y, though this could change in 500 years.' One would expect more of a realist stance towards explanation.

contextualist account, the problem is handled in stride: theories succeed because the sought-after result is found, not because the theory is true.[134]

But the following problem presents itself. Call a theory qualitatively successful if it accurately predicts some phenomenon *p* rather than not predicting *p*. And call a theory quantitatively successful if the theory accurately predicts the degree to which *p* occurs rather than not predicting accurately the degree to which it occurs. We can say that the contrast class for the qualitative success of Txy (the theory that x occurs in degree y) is 'p, q...n' (where p, q...n are not equal to x) and the contrast class for the quantitative success of Txy is '*xr*, *xs*,...*xn*'' (where r, s...n are nor equal to y). So Einstein and de Haas enjoyed quantitative success in that their theory accurately predicted that *g* would be equal to 1. They enjoyed qualitative success in that they correctly predicted that the iron rod would turn.

The problem is that contextualism accounts for quantitative success but not qualitative success. As John Greenwood puts it, "there was no guarantee that they would get the result that conformed with their theoretical expectations. ... [I]f Einstein and de Haas had not had their theory, the experimental results would not have counted as a success- but this does nothing to explain their ability to predict it successfully given their theory." Alternatively, as Michael Levin pithily puts it, contextualism doesn't explain "why any of the experiments got anywhere near the predicted value. Why did even one succeed?"

[134] This raises again the issue of facticity. As I have suggested throughout this essay, we can avoid this issue by talking about false but apparently successful theories and casting the problem of success as the problem of explaining how it is that false theories enjoy such apparent successes. This strategy resembles the retroactive realist strategy I dismiss in the previous note. But the contextualist is happy to index

The problem is compelling: Given their theoretic predilections, Einstein and de Haas might have been inclined to reject certain results (on which g was not equal to 1) and accept other results (on which g was equal to 1). The rate at which the iron rod turned was, we can grant, subject to interpretative difference: as noted above, movements that fell significantly out of the expected range were dismissed as the results of systematic errors, thereby allowing theoretic dispositions to affect the determination of angular velocity. But the fact that the rod turned is not subject to such interpretative difference. What enabled Einstein and de Haas to correctly predict that the rod would turn? Contextualism does not seem to possess the resources to answer this question.

The contextualist might respond by claiming that scientists observe regularities in the world, and they develop theories to explain them, often in terms of unobservable entities and processes. On this line of argument, the experiments got near to the predicted values because there is a real world out there, and there are real world constraints on what happens when iron bars get magnetized. The facts are a given, the contextualist might say, known in the same way that I know that when I put bread in my toaster and push down the lever, the bread will be lightly cooked: I know this not because I know anything about heating coils, but because I know, based on experience and the reports of others, that this regularity obtains.[135] Similarly, Einstein and de Haas identified a particular

(contextualize) success – so this is not an unhappy result – while the realist wants a more robust theory-world connection that seems to be precluded by this strategy.

[135] This talk of empirical regularities makes contextualism sound like van Fraassen's constructive empiricism: the scientist latches onto an empirical regularity and explains that regularity in theoretic terms, but this does not compel belief in the entities postulated by that explanation. What the views have in common, of course, is the commitment to empirical regularities. But the contextualist goes beyond the constructive empiricist in that he offers an analysis of predictive success, whereas van Fraassen fails even to account for explanatory success. The views share a commitment to a basic empirical regularity, but they put their commitments to very different uses.

regularity – the movement of iron bars when brought into proximity with magnets – and they sought to explain this regularity in terms of a particular ontology. This explains the qualitative success of their theory, and contextualism explains its quantitative success: their account of the regularity succeeded not because their theory accurately modeled the world, but because their theoretic expectations of what did the explaining influenced their experimental inquiries into what did the explaining.

But this just pushes the problem of qualitative success back a step. The theory said the rod would turn and the rod turned. How did Einstein and de Haas know this? Because it was an empirical regularity and scientists are in the business of identifying empirical regularities and offering accounts of them. Fine- but how is it that scientists are able to capture empirical regularities at all? Why is it that scientific theories – rather than, say, blind guesses or tea leaves – regularly capture these regularities? The contextualist is utterly unable to answer this question – it remains a brute fact about science – and so leaves the contextualist unable to explain why scientific theories yield accurate predictions at all.

That the contextualist can account for quantitative success is, I suppose, worth something. But it's not worth much. For one thing, the account of quantitative success is parasitic upon the account of qualitative success: the contextualist is saying that, for any successful theory, he can tell you why that theory succeeds to the degree to which it does. Ask him why the theory succeeds at all, though, and he is silent. Contextualism thus offers an incomplete picture of scientific practice. Worse, the picture it does offer presupposes that science is successful—but this is precisely what needs explaining![136]

6.4 a second problem for contextualism: failed theories

Another problem for contextualism is that it seems to preclude the possibility of theoretic failure.[137] Consider the following: Michelson and Morley, theoretically disposed to believe in the existence of the aether, conducted experiments to measure the speed at which the earth travels in the aether. Their experiment revealed none of the results that should have obtained if the aether theory were true. If the contextualist were right, the experimenters would have made the results fit with their expectations; they would not have acquiesced in the face of recalcitrant data. Indeed, if the contextualist were right, there would be no recalcitrant data; Michelson and Morley would have thrown out bad measurements and interpreted their data in such a way as to vindicate their theory. It would seem that this serves as a counter-example to the contextualist account of success.

Again, the problem poses a serious threat to contextualism: we want our philosophies (and sociologies) of science to accurately reflect scientific practice, and one feature of scientific practice is that theories sometimes fail to yield correct predictions. If experimental evidence is always open to interpretation, though, and if predictive success is obtained by the selective interpretation of data, it would seem to follow that predictive failures never occur.

[136] Indeed, the problem with which most of the writers on the subject are concerned is the ability of science to repeatedly 'get it right': why, they ask, are scientific theories able so often to capture empirical regularities. An account of the success of science that presupposes that science has this ability is, then, doubly inadequate.

The contextualist might respond by appealing to the distinction, introduced earlier, between qualitative success and quantitative success. That is, the contextualist could argue that contextualism explains only why theories are quantitatively successful and contend that the sort of failure at issue here is a failure of qualitative success—something that contextualism cannot explain but does allow for. But this would be a rather hollow victory. We want to know why theories sometimes fail to yield accurate predictions. The present suggestion is that such failures are qualitative failures, and contextualism does not have any truck with qualitative success or failure. But this doesn't explain why theories fail; it just acknowledges that they do. Nor does it identify any feature of scientific practice – theoretic or otherwise – in virtue of which theories fail. In short, the contextualist tells us nothing about scientific failure, other than that it happens. We deserve more from a theory of science.

6.4.1 another response to the problem of failed theories

The better strategy for the contextualist might be to appeal to the sociological (controversy studies) tradition out of which the theory develops. On this line of argument, the contextualist would argue that theory does inform experiment and interpretation of experimental results, but this occurs over time: if the first experiment fails to yield the right results, new instruments are built, theoretic disputes are entered into, experiments are adjusted, etc. This all takes time, and this makes room for scientific ontologies and knowledge to be constructed: over the course of attempts to fit theory and data together, theoretic contexts sometimes change and, as Pickering suggests, new theories are introduced that explain the recalcitrant data and offer new theoretic directions

[137] Thanks to Michael Levin for pushing me on this point.

in which to travel. This is why theories sometimes fail- theoretic contexts change before the theories can succeed.

Consider again the Michelson-Morley experiments.[138] We know that, in his first aether drift experiment of 1881, Michelson paid exquisite attention to the experimental set-up and thus left little room for immediate adjustment of instrumentation or results. The results did not match up with expectations, and no obvious adjustments presented themselves. Lorentz suggested that the null result might be due to a failure to account for the effect of the wind on the experimental apparatus, resulting in an aether drift measurement that was hidden by 'experimental noise'. This led Michelson to adjust the apparatus and, with Morley, try the experiment again—six years later. When this experiment also produced a null result – that is, when the results of this second experiment did not immediately lend themselves to a favorable interpretation – Michelson simply lost interest in the issue and moved on: "Michelson seems to have been so disappointed at the result that instead of continuing he immediately set about working on a different problem..." (CP 37). What's worth noticing is that, "in spite of Michelson's own lack of interest in his findings, discussion did not cease. The results were seen as a 'cloud' in the otherwise clear sky of physics. Numerous explanations were put forward in an attempt to show how the existence of an aether was compatible with the null results" (CP 38). Morley continued to work on the problem for the next 20 years; in the early 1920's, Dayton Miller even conducted a series of experiments that showed a positive, non-null result.[139]

[138] I here follow Collins and Pinch (1993: 27-43), hereafter abbreviated 'CP'.
[139] Indeed, as Arnold Koslow reminds me, even Einstein didn't think the aether theory had been eliminated.

Over the course of these almost fifty years of experimentation, however, Einstein's theory of relativity came to be accepted and, with its acceptance, the need for an explanation of the null result waned. So the facts did not immediately lend themselves to interpretation by the lights of theory and, while scientists were making experimental and theoretic adjustments, another theory came along, captured the scientific community, and changed the proverbial rules of the game: aether drift theories were now considered false, so any non-null result would not be taken seriously: "the sheer momentum of the new way in which physics was done – the culture of life in the physics community – meant that Miller's experimental results were irrelevant" (CP 42). Hence, by the time Miller finally achieved Michelson's goal – when he finally fit theory and fact together – it was too late: no one was listening.

The case study shows a long history of attempts to fit data and theory. While scientists were attempting to make this happen, another theory came along that explained the results and provided a context in which the null result made sense. At this point, the question of whether the aether drift theory was right became moot. Michelson and Morley didn't accept failure; rather, another theory came along to displace the aether drift theory before they could make that theory and the data coincide. Michelson ceased working on the problem for straightforward biographical reasons: he was dejected by the failure of the results to immediately gel with his theory and was not interested in pursuing the matter any longer.[140] Others continued working on the problem (and eventually

[140] See Collins and Pinch (27-43, esp. 37-38). The sociologist might pursue the matter in something like the following way: Michelson had been working on this problem for at least ten years. He did not want his career to consist of work on this one issue; he did not want to be a 'one-trick pony'. So he moved on before achieving success, but for workaday reasons. Did he have tenure? Was he satisfactorily situated in

succeeded in getting the results they wanted). Theory guided the initial experimental set-up, leading to a second experiment and a series of additional experiments, all of which sought to introduce adjustments to the experiment, theory, and instruments such that the theory of aether drift would be substantiated. While this process was underway, the theory of relativity displaced the theory of aether drift and, in so doing, displaced the matter of the success of the aether drift experiments. "In spite of [Miller's successful 1924 experiment] the argument in physics was over. Other tests of relativity … indirectly bolstered the idea that the theory of relativity was correct …" (CP 42).

In short, the contextualist can argue that the aether-drift theory failed because, by the time scientists were able to fit the data and the theory, the social climate of professional science had changed. More generally, theoretic (qualitative) failures are the result of shifts of theoretic climate: it takes time for theory and data to be fit together, and while such fits are being worked out, new theories come into play and offer solutions to the problems. These alternative solutions appear to be successes for the new theories and seem to show that the old theories failed, but in fact they just foreclose on the need to work on the old theories. The old theories aren't shown to be false; just outdated.

6.4.2 problems with the case study solution

This is, I say, the best response the contextualist can offer to the problem of theoretic failure: it is consistent with the sociological tradition out of which contextualism is developed and it is stoutly non-realist in it's insistence that the source of failure, like

---

the field? If not, there were strong reasons for him to shift his professional attention to problems that could be solved more quickly.

success, lies beyond the resources of the theory itself. But, even if the response works for this case, it is quite a leap to go from one case study to all of science; more work would need to be done to show that other instances of theoretic failure can be accounted for in sociological terms. Moreover, a host of questions plague this sort of case-by-case solution: how many case studies would need to be done? from how many fields of science? how much detail would each case study need to involve? If the contextualist plans to account for the fact that theories sometimes fail by examining particular instances of failure and basing a generalization about scientific practice on them, he has to give us a sense of what will count as an adequate generalization, and it seems reasonable that to require, at a minimum, answers to questions like these.

Fortunately, we needn't answer these questions; there are problems with the contextualist account of failure that keep it from getting off the ground.

One problem is that it is likely that the theoretical context changed *because* of the null result: Michelson and Morley's failure to substantiate aether drift theory led to a crisis in the scientific community, and this crisis led to the adoption of another theory. If this is the case, then it is disingenuous, at best, for the contextualist to claim that we can explain Michelson and Morley's failure by appeal to the adoption of a new theory before Michelson and Morley could make the data and theory fit. If it was their failure to effect such a fit that led to the adoption of the new theory, then we cannot reasonably explain that failure by appeal to the adoption of a new theory.

The contextualist is likely to respond as follows: part and parcel of the sociological view of theory development is the recognition that theories aren't accepted all at once or immediately. Just as work on the aether drift theory continued for a long time, so did work on the theory of relativity: "It should not be thought … that Einstein's ideas were uniformly accepted upon their publication. The battle lasted for several decades." (CP 39). Over the course of these years, the theory of relativity and the aether drift theory were prodded and tested on many fronts. There were no crucial experiments; the theory of relativity was accepted for many reasons, all of them but one independent of the Michelson-Morley experiment. As Einstein's theory came to be accepted, it "gained ground by explaining the Michelson-Morley anomaly. Because relativity was strong, it seemed the natural template through which to interpret the 1919 observations. Because those observations then supported relativity further, the template was still more constraining when it came to dealing with Miller's 1925 observations" (CP 52).

Even if we grant that there are no crucial experiments and that the theory of relativity was accepted for reasons other than (but including) the null result, though, there remain problems with the case study solution. All it does is posit that theories fail and show how these failures are followed by the introduction of alternative theories that succeed. Why were Michelson and Morley unable to interpret the data by the lights of their theory? Because the experiments needed to be altered in order to bring the results more in line with theoretic expectations and, while those alterations were being made, the theory of relativity came along and was accepted. When theory and data finally were brought into alignment, no one was interested. Okay. But all this comes to is the claim that theory and data could not be fit together, and then they could. There remains the question of why they could not manage a fit of theory and data.

The contextualist tells us that they could, but it took time—so much time, in fact, that it fell to Dayton Miller, working almost 50 years later, to finish the project they started. What this asks us to believe, though, is that Miller's experiment provides sufficient evidence for aether drift theory but is ignored only because scientists are already committed to the theory of relativity. As Michael Levin suggests, this is rather implausible; the experimental history of null results provides good reason to think that there simply was no aether drift to measure. It is more likely that Miller's experiment was flawed than it is that the long history of experimental failure was mistaken.

This gives rise to a serious methodological problem. In order to determine whether Miller's experiment was well-designed and executed, we would need to attempt to replicate it. But the experimenter's regress would re-assert itself: the new experiment would either be informed by the theory of relativity or by the aether drift theory.[141] This suggests that an unbiased assessment of Miller's experiment is not available, and this would mean that the case study account of Michelson and Morley's failure – and, more generally, the case study approach to the problem of theoretic failure – is every bit as fallible as the initial experiments.

[141] Or by some other theory. The sociologist, recall, would contend that this kind of prior sorting of possibilities is one of the reasons experiments yield the results that are expected. (See note 132).

We started out by noticing that contextualism seems to preclude the possibility of failure: data should always be interpreted in a theoretically favorable way. The contextualist response to this challenge is that it sometimes takes time for interpretations to be worked out and, during that time, other theories are accepted, and this makes it seem as though the initial theory failed. But this response does little more than grant the initial problem-it admits that data is not always interpreted by the lights of theory. The initial theory did fail, and the contextualist's response does not explain why this is the case. To be sure, it allows for it, but only in the sense that it admits that failure occurs. What is needed is an account of how the theoretical resources of contextualism make room for such failure, and the contextualist has nothing to say on this score that does already presuppose that such failures occur.

6.5 the implications of these problems for the contextualist account of success

The contextualist is unable to account for the fact that theories fail. While he might allow that theories fail, his theory lacks the resources to account for such failure. At the end of the day, the contextualist has as little to say about the problem of failure as he has to say about the problem of qualitative success: it happens, but he can't explain why. Where does this leave us?

We set out to see whether the sociology of science (specifically, constructivism) could succeed where the philosophy of science failed. We determined that the only version of constructivism that could work was ontological constructivism – the view that ontological commitments stem from decisions about which theoretical constructs to take as real and are the results of historical processes – and we decided that Pickering's opportunism in context offered the most promising account of how this historical process works: in the space left by the fallibility of experimental outcomes, theoretic concerns influence the direction of experiments and so influence experimental outcomes, including the determination of what exists. Contextualism is intended to stretch Pickering's theory, which is an account of theory choice, to cover the predictive success of science.

What we found was that the theory cannot be stretched that far. Contextualism allows us to make sense of a theory's ability to predict the *degree to which* something will occur, but it has nothing to say about how the theory is able to issue in a correct prediction *in the first place*. If we take there to be two kinds of success – the success of science to 'get it right' at all, or qualitative success, and the success of particular theories to 'get the details right', or quantitative success –contextualism allows us to explain only the latter kind of success. As for the former kind, we seem forced to conclude, with Popper, that there is no explanation of success to be had. This is serious enough to compromise the meta-theoretical integrity of contextualism: the most basic fact of science's success (its ability to tell us what's going to happen next) is not only assumed, but must be assumed in order for the contextualist's account of quantitative success to get going.

The contextualist might respond by noting that, once we accept the anti-realist critique of realism, the realist also is without an explanation of science's ability to explain the qualitative success of science, and this places realism and contextualism on all fours, vis-à-vis qualitative success, and it leaves contextualism with a comparative advantage vis-à-vis quantitative success. But let's not kid ourselves. If contextualism's victory over realism isn't hollow, it is certainly lacking. What we want from a theory of science (be it philosophical or sociological) is a full-blown account of how science works. Insofar as one major feature of science is its ability to 'get it right', contextualism fails to serve as an adequate account of science.

Any glimmer of an advantage for contextualism is extinguished by the fact that contextualism lacks the resources to explain why theories fail. If realism does nothing else, it offers us a clear account of why theories fail to issue in accurate predictions: the world does not work the way the theories say it does. This resource is not available to the contextualist, who is forced to adopt a case study approach to the issue of failure, according to which failure amounts to an inability to interpret the data by the lights of the theory before another theory that can explain the data is accepted. The contextualist claims that examination of the details of cases of theoretic failure shows that this is what happens and so concludes that contextualism can account for the fact that theories sometimes fail. But this response is inadequate: even if we assume that it can be applied to more than one historical episode (and even if we ignore the problem of determining adequate criteria of applicability), it fails to identify any feature of scientific practice in virtue of which the data cannot be interpreted by the lights of the theory; all it does it assert that this is the case. Even if it were the case, the theory asks us to believe that long histories of experimental failure are not reasons to believe the theories are false and that case studies reveal that the abandoned theories are eventually vindicated (but not taken seriously because of the change of theoretic climate). This is intuitively difficult, to say the least. Moreover, and more seriously, the vindicated theories would be judged from particular theoretic standpoints and the problem of interpretation would re-assert itself.

The initial appeal of contextualism was that it promised to account for the apparent successes of false theories- something that realism cannot do. What we have found is that contextualism allows us to understand why false theories can enjoy predictive success, but it doesn't allow us to understand why true theories succeed or fail. Realism allows us to understand why true theories succeed and fail, but it doesn't allow us to understand why false theories appear to succeed.[142] This is an odd result. Perhaps there is some lesson to be learned from it. If there is, I suspect it will be one the sociologists do not want to hear: before Bloor, sociology was invoked when theories failed but not when they succeeded; the sociology of science – specifically, the symmetry postulate – was developed to correct this. Bloor wanted sociology to be invoked in the explanation of truth as well as falsity. But here we see the asymmetry re-presenting itself: sociology works for false theories and realism works for true theories. Indeed, our sociological investigation of success seems to offer a principled reason to reject the sociologist's symmetry postulate!

In any event, it is clear that contextualism cannot account for the success of science. While this is not a knock-down argument against sociological approaches to success – there might be another version of constructivism that fares better – it is sobering. Our chapter five analysis of the sociology of science suggested that the only sociological

[142] Thanks to Arnold Koslow for pointing out this peculiar symmetry. More generally, thanks to Koslow for guiding me through and to the end of this.

theory that stands a chance of explaining success is ontological constructivism. Pickering's opportunism in context offers the best statement of that view, and contextualism is an extension of Pickering's theory. Absent the offer of an alternative version of ontological constructivism, we can conclude that sociological accounts of success offer no more hope than do philosophic accounts.

The reader will recall that realism, while intuitively appealing, was found to have little philosophic merit. Anti-realism was found to have just as little merit and, indeed, was shown to collapse into realism. So where are we? Most people's intuitions scream realism. Mine do. Airplanes fly because we know a thing or two about the laws of aerodynamics. Lamps work because we know about the laws of electricity. The problem – and I submit that this is the right way to understand the problem of success for science – is that our intuitions tell us realism is true but the arguments tell us it is false. As we noted at the outset of this investigation, however, intuitions are not arguments, and things are not always as they appear. What we wanted was a philosophical account of the success of science that would either substantiate our intuitions or give lie to them. We found neither; the arguments support Popper's conclusion that success is inexplicable. At the end of the day, all we are left with is this and our realist intuitions.

Bibliography

Abela, P. 1996. Is Less Always More? An Argument Against the Natural Ontological Attitude. *The Philosophical Quarterly*: 72-76.

Barnes, B., D. Bloor, and J. Henry 1996. *Scientific Knowledge*. Chicago: University of Chicago Press.

Boyd, R. 1973. Realism, Underdetermination, and a Causal Theory of Evidence. *Nous*: 1-12.

----------. 1981. Scientific Realism and Naturalistic Epistemology. in P. Asquith and R. Giere (eds.) *PSA 1980*, East Lansing: Philosophy of Science Association.

----------. 1984. The Current Status of Scientific Realism in Leplin 1984, 41-82.

----------. 1985. The Logician's Dilemma: Deductive Logic, Inductive Inference and Logical Empiricism. *Erkenntnis 22*: 197-252.

----------. 1989. What Realism Implies and What it Does Not. *Dialectica 43*: 5-29.

----------. 1990. Realism, Conventionality, and 'Realism About'. in G. Boolos (ed.) *Meaning and Method: Essays in Honor of Hilary Putnam*, Cambridge: Cambridge University Press.

----------. 1991. Realism, Anti-Foundationalism, and the Enthusiam for Natural Kinds. *Philosophical Studies 61*: 127-148.

Brandon, E.P. 1997. California Unnatural: On Fine's Natural Ontological Atitude. *The Philosophical Quarterly*: 232-235.

Bunzl, M. 1994. Discussion: Scientific Abstraction and the Realist Impulse. *Philosophy of Science 61*: 449-456.

Collins, H.M. 1992. *Changing Order*. Chicago: University of Chicago Press.

Collins H.M. and T. Pinch. 1993. *The Golem: What Everyone Should Know about Science*. Cambridge: Cambridge University Press.

Collins, R. 1994. Against the Epistemic Value of Prediction Over Accomodation. *Nous 28*: 210-224.

Devitt, M. 1991a. *Realism and Truth, 2nd ed.* Princeton: Princeton University Press.

----------. 1991b. Aberrations of the Realism Debate. *Philosophical Studies 61*: 43-63

----------. 1991c. Realism Without Representation: A Response to Appiah. *Philosophical Studies* 61: 75-77.

Dummett, M. 1982. Realism. *Synthese 52*: 55-112.

----------. 1993. *The Seas of Language*. Oxford: Oxford University Press.

Earman, J. and C. Glymour. 1980. Relativity and Eclipses: The British Eclipse Expeditions of

1919 and their Predecessors. *Historical Studies in the Physical Sciences 11*. 49-85.

Ellis, B. 1985. What Science Aims to Do. in P. Churchland and C. Hooker (eds.) *Images of Science: Essays on Realism and Empiricism*. Chicago: University of Chicago Press.

Fine, A. 1984. The Natural Ontological Attidude. in Leplin 1984, 83-107.

----------. 1986. Unnatural Attitudes: Realist and Instrumentalist Attachments to Science. *Mind 95*: 149-179.

----------. 1989. Truthmongering: Less is True. *Canadian Journal of Philosophy 19*: 611-616.

----------. 1991. Piecemeal Realism. *Philosophical Studies 61*: 79-96.

----------. 1996a. *The Shaky Game*. Chicago: University of Chicago Press.

----------. 1996b. Science Made Up. in Gallison. P. and D.Stump (eds.) *The Disunity of Science*. California: Stanford University Press.

Franklin, A. 1986. *The Neglect of Experiment*. Cambridge: Cambridge University Press.

Friedman, M 1979. Truth and Confirmation. *Journal of Philosophy 76*: 361-382.

Galison, P. 1987. *How Experiments End*. Chicago: University of Chicago Press.

Golinski, J. 1998. *Making Natural Knowledge*. Cambridge: Cambridge University Press.

Goodman, N. 1983. *Fact, Fiction and Forecast*. Cambridge: Harvard University Press.

----------. 1984. *Of Mind and Other Matters*. Cambridge: Harvard University Press.

Gutting, G. 1980. Scientific Realism vs Constructive Empiricism: A Dialogue. *Monist 65*: 336-349.

Hacking, I. 1983. *Representing and Intervening*. New York: Cambridge University Press.

---------------1999. *The Social Construction of What?* Cambridge: Harvard University Press.

Hardin, C. and A. Rosenberg. 1982. In Defence of Convergent Realism. *Philosophy of Science 49*: 604-615.

Hellman, G. 1983. Realist Principles. *Philosophy of Science 50*: 227-249.

Hooker, C. 1987. *A Realistic Theory of Science*. SUNY Press.

Kitcher, P. 1993. *The Advancement of Science*. Oxford University Press.

Klee, R., ed. 1999. *Scientific Inquiry: Readings in the Philosophy of Science*. New York: Oxford University Press.

Knezevich, L. 1989. Truthmongering: An Exercise. *Canadian Journal of Philosophy 19*:

603-610.

Kukla, A. 2000. *Social Constructivism and the Philosophy of Science*. London: Routledge.

Ladyman, J., I. Douven, L. Horsten, and B. van Fraassen. 1997. A Defence of van Fraassen's Critique of Abductive Inference: Reply to Psillos. *The Philosophical Quarterly 47*: 305-321.

Latour, N. and S. Woolgar. 1986. *Laboratory Life: The Social Construction of Social Facts*. Princeton: Princeton University Press.

Laudan, L. 1984a. Explaining the Success of Science: Beyond Epistemic Realism and Relativism. in A. Tauber (ed.) *Science and the Quest for Reality*. New York: New York University Press.

----------. 1984b. Discussion: Realism without the Real. *Philosophy of Science 51*: 156-162.

----------. 1984c. A Confutation of Convergent Realism. in Leplin 1984, pp. 218-249.

Levin, M. 1984. What Kind of Explanation is Truth? in Leplin 1984, pp. 124-139.

Levin, M. 1990. Realisms. *Synthese 85*: 115-138.

Levin, Ma. 2000. Upholding Truth: Objectivity vs. Skepticism and Nihilism. In Louis Pojman (ed.) *Classics of Philosophy: the 20th Century*, 2nd ed. Oxford: Oxford Univrtsity Press.

Leplin, J., ed. 1984. *Scientific Realism*. Berkeley: University of California Press.

----------. 1997. *A Novel Defense of Scientific Realism*. New York: Oxford University Press.

Lipton, P. 1991. *Inference to the Best Explanation*. New York: Routledge.

McMichael, A. 1985. van Fraassen's Instrumentalism. *British Journal for the Philosophy of Science 36*: 257-272.

McMullin, E. 1984. A Case for Scientific Realism. in Leplin 1984, pp. 8-40.

Musgrave, A. 1989. NOA's Ark- Fine for Realism. *The Philosophical Quarterly 39*: 383-398.

Nanda, M. 1997. Against Social De(con)struction of Science: Cautionary Tales from the Third World. *Monthly Review 48*: 1-20.

Pappineau, D. (ed.). 1996. *The Philosophy of Science*. Cambridge: Oxford University Press.

Pickering, A. 1984. *Constructing Quarks*. Chicago: University Press.

Psillos, S. 1996. *Scientific Realism: How Science Tracks Truth*. London: Routledge.

----------. 1997. How Not to Defend Constructive Empiricism: A Rejoinder. *The Philosophical Quarterly 47*: 369-372.

Putnam, H. 1975. *Mathematics, Matter, and Method.* Cambridge: University Press.

----------. 1984. What is Realism? in Leplin 1984, 140-153.

Rouse, J. 1991. The Politics of Postmoden Science. *Philosophy of Science 58*: 607-627.

Schlagel, R. 1991. Critical Notice: Fine's "Shaky Game" (And Why NOA is no Ark for Science). *Philosophy of Science 58*: 307-323.

Smart, J.J.C. 1963. *Philosophy and Scientific Realism.* London: Routledge.

van Fraassen, B. 1980a. *The Scientific Image.* Oxford: Clarendon Press

----------. 1980b. Theory Construction and Experiment: An Empiricist View. in P. Asquith and R. Giere (eds.) *PSA 1980*, East Lansing: Philosophy of Science Association.

----------. 1985. Empiricism in the Philosophy of Science. in P. Churchland and C. Hooker (eds.) *Images of Science: Essays on Realism and Empiricism*, Chicago: University of Chicago Press.

----------. 1989. *Laws and Symetry.* Oxford: Clarendon.

----------. 1991. The Pragmatics of Explanation. In R. Boyd, P. Gasper, and J.D. Trout (eds.) *The Philosophy of Science*, Cambridge: MIT Press.

Wartofsky, M. 1991. How To Be a Good Realist. in G. Munewar (ed.) *Beyond Reason*, the Netherlands: Kluwer Academic Publishers.

Worral, J. 1984. An Unreal Image. *British Journal for the Philosophy of Science 35*: 65-80.

------------------ 1996. in Papineau 1996, 139-165.