

Durham E-Theses

*How and Why to Localise the Scientific Realism
Debate: Making Historical Arguments in the
Scientific Realism Debate Compatible with
Methodological Pluralism*

ALEXANDER JOSHUA ELLIOTT

How to cite:

ELLIOTT, ALEXANDER JOSHUA (2023) How and Why to Localise the Scientific Realism Debate: Making Historical Arguments in the Scientific Realism Debate Compatible with Methodological Pluralism. Masters thesis, Durham University.

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a <https://etheses.durham.ac.uk/id/eprint/15253/> is made to the metadata record in Durham E-Theses
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full Durham E-Theses policy](#) for further details.

Abstract:

I argue for a new way of localising historical arguments in the scientific realism debate. We should see historical arguments as attempts to empirically assess scientific methods. For such arguments to be good, they need to be about a single method. Therefore, only if there is a unified method of science can historical inductions on science license general conclusions about the epistemic status of current science. However the consensus seems to be that there is no such unified scientific method. Various versions of methodological pluralism seem to undermine any attempt to assess scientific methods through historical means, as they make it hard to see methods as persisting through theory change or as applying beyond a very specific field. In particular, views of scientific methods that see them as highly context specific seem to undermine any kind of historical realism debate.

I attempt to outline a way in which we can individuate scientific methods in order to empirically test them. I also argue that the impact of context can be accounted for in a way that still leaves room for historical assessments of methods if we categorise contexts according to types of difficulty. The view of the historical scientific realism debate we end up with is one in which various methodological resources are argued to be either unreliable or reliable for a given type of difficulty, based on evidence from the history of science. These conclusions about the reliability of methods may be relevant to the epistemic status of a given theory, but establishing which methods and difficulties are present in an actual scientific context requires detailed engagement with local evidence. I compare my position with other localist views and explain how it offers more role for historical inductions on science than some other localist writers.

How and Why to Localise the Scientific Realism Debate: Making Historical Arguments in the
Scientific Realism Debate Compatible with Methodological Pluralism

Alexander Joshua Elliott

Submitted for an MPhil in Philosophy

Department of Philosophy

Durham University

2023

Contents

1. Chapter 1: Introduction (p.7)

- a. Section 1: the Debate Over Scientific Realism
 - i. Section 1.1: The success to truth inference (STI)
 - ii. Section 1.2: Anti-realism and the Pessimistic Induction (PI)
 - iii. Section 1.3: Realist defence 1: Apparently successful-but-false theories were not really false
 - iv. Section 1.4: Realist defence 2: Apparently successful-but-false theories were not really successful
 - 1. Section 1.4.1: Novel Predictive Success
- b. Section 2: Summary of Thesis Argument by Chapter

2. Chapter 2: Can the Predictive Requirement on the success to truth inference be Independently Motivated? (p.17)

- a. Section 1: Why independent motivation is needed
- b. Section 2: What would it require for us to say that there was independent motivation for specifying that 'success' means predictive success?
- c. Section 3: Is the predictive success inference independent from the non-predictive success inference?
- d. Section 4: Why prediction and explanation are related concepts
- e. Section 5: Theory choice when we have to balance predictive and non-predictive success
- f. Section 6: Independent grounds for thinking predictive success is epistemically superior
- g. Section 7: How much does the move to a predictive success to truth inference really need independent motivation?

3. Chapter 3: Case studies in non-predictive success: history of evolutionary biology 1859-1937 (p.24)

- a. Historical introduction
- b. Section 1: the Case for Evolution in the Nineteenth Century
 - i. Section 1.1: Evolution as the only scientific option
 - ii. Section 1.2: The explanatory power of evolution
 - iii. Section 1.3: Intermediary fossils
 - iv. Section 1.4: Discussion
- c. Section 2: Natural Selection

- i. Section 2.1: Why was natural selection not widely accepted in the nineteenth century?
 - ii. Section 2.2: What led to natural selection becoming widely accepted?
 - iii. Section 2.3: What does the case for natural selection tell us about prediction vs. explanation?
 - d. Conclusion
- 4. **Chapter 4: A Different Approach to the Realism Debate (p.37)**
 - a. Introduction
 - b. Section 1: the Realism Debate as a Methodological Debate
 - c. Section 2: Problems for Testing Methodological Principles Empirically
 - d. Section 3: Can 'Truth' Feature in a Testable Methodological Principle
 - e. Section 4: Implications of Seeing the Realism Debate as a Methodological Debate
 - i. Section 4.1: Two Possible Errors for Historical Assessments of Methods
 - 1. Section 4.1.1: Conflating Topics
 - 2. Section 4.1.2: Conflating Methods
 - 3. Section 4.1.3: What these Types of Errors tell us about the STI and the PI
 - ii. Section 4.2: How well does the STI Reflect Actual Scientific Methods?
 - iii. Section 4.2: Methodological Universalism, Methodological Pluralism, and their Implications for the Realism Debate
- 5. **Chapter 5: Methodological Pluralism and Individuating Methods for Historical Assessment (p.52)**
 - a. Introduction
 - b. Section 1: Universalism and its problems
 - c. Section 2: Two Routes to Pluralism
 - i. Section 2.1: Multiple Sets of Rules
 - 1. Section 2.1.1: Kuhn
 - 2. Section 2.1.2: Hacking
 - 3. Section 2.1.3: Discussion of Multiple-Systems-of-Rules Type Methodological Pluralism
 - ii. Section 2.2: Contextuality and Defeasibility of Rules
 - 1. Section 2.2.1: Feyerabend
 - 2. Section 2.2.2: Norton
 - 3. Section 2.2.3: Maddy and Fine

- 4. Section 2.2.4: Saatsi on Form and Content Driven Arguments
- d. Section 3: Lessons of Methodological Pluralism for the Scientific Realism Debate
- 6. **Chapter 6: Localism and Individualism: Choosing the Right Level for Evaluating Historical Arguments in the Scientific Realism Debate (p.73)**
 - a. Introduction
 - b. Section 1: Arguments for Localism
 - i. Section 1.1: Asay
 - ii. Section 1.2: Park's Objections to Asay
 - iii. Section 1.3: Fitzpatrick
 - iv. Section 1.4: Henderson's Objections to Localism
 - c. Section 2: How the Historically Motivated Anti-Realist can Still Attack Localism
- 7. **Chapter 7: Individuating Topics in Historical Assessments of Methods (p.88)**
 - a. Introduction
 - b. Section 1: How can the Anti-Realist Characterise the Limits of their Scepticism?
 - c. Section 2: Types of Difficulty in Science
 - i. Section 2.1: Inaccessibility
 - ii. Section 2.2: Unfamiliarity
 - iii. Section 2.3: Conceptual Difficulties
 - iv. Section 2.4: Inconstant and Interactive Subject Matter
 - v. Section 2.5: Types of Question
 - vi. Section 2.6: Analysis of Variety of Difficulties Found in Science
 - d. Section 3: Relativising Historical Assessments of Methods to Types of Difficulty
 - i. Section 3.1: Limiting the Scope of Sceptical Arguments
 - ii. Section 3.2: Advantages of Relativising Historical Assessments of Methods to Types of Difficulty
 - iii. Conclusion
- 8. **Chapter 8: Review, Comparison with Other Work in the Field, and Conclusion (p.106)**
 - i. Introduction
 - ii. Section 1: Summary of Thesis Argument
 - iii. Section 2: Similarities and Differences between my Position and Other Authors
 - 1. Section 2.1: Saatsi on Replacing Recipe Realism
 - 2. Section 2.2: Fine, Maddy, Park and Vickers
 - 3. Section 2.3: Asay

4. Section 2.4: Henderson
 5. Section 2.5: Summary of how my work first into the debate
 - iv. Section 3: Suggestions for Further Work
9. **References (p.114)**

List of Abbreviations

- **NMA:** No Miracles Argument
- **NOA:** Natural ontological attitude
- **NPS:** Novel Predictive Success
- **NPSTI:** Non-Predictive Success to Truth Inference
- **PI:** the Pessimistic Induction
- **PSTI:** Predictive Success to Truth Inference
- **STI:** the Success to Truth Inference

Statement of Copyright

The copyright of this thesis rests with the author. No quotation from it should be published without the author's prior written consent and information derived from it should be acknowledged.

Chapter 1: Introduction

Section 1: the debate over scientific realism

The debate over scientific realism is a debate over the extent to which we are justified in believing that our best current scientific theories provide literally true descriptions of a reality that is independent of our thoughts about it. Scientific realists argue that we are justified in believing the literal truth of at least some of current science. Scientific anti-realists dispute at least some of the knowledge claims of scientific realists.

Already, characterising the scientific realism debate has required a considerable degree of vagueness. Most of this has derived from the need to specify what the objects of realists' and anti-realists' disputes are. I have characterised the thing they differ on as the epistemic status of 'science' or 'our best current science'. A theme I will repeatedly be returning to later is how diverse these categories are. It might immediately strike the reader that they might want to be realists about some areas of science and anti-realists about others. That is a situation that most philosophers are in regardless of whether they identify as realist or anti-realist (Vickers 2022, Stanford 2021), despite the tendency of both camps to talk about 'realism' and 'anti-realism' as though they were global attitudes towards science.

To add clarity to what the scientific realism debate is about it will be helpful to introduce different levels at which one can be realist or anti-realist. On the one hand, there is the global level. Here, we either regard science in general as mostly providing claims that we are in a position to know are literal truths, or mostly failing to do so. A philosopher holding any view at this level can be described as having a 'global' approach or as being a 'globalist' (Asay 2017). Moving down, we could perhaps be a realist or anti-realist at the level of whole scientific fields of research. For example, we might say that physics generally provides things we can know to be literal truths, but that economics mostly fails to do so. Moving down further, we could hold a scientific realist or anti-realist attitude towards individual theories. For example, maybe realism is justified about the theory of evolution by natural selection, but not justified about the theory of general relativity. Finally, we might be realists or anti-realists about individual scientific claims such as 'the mass of the earth is $5.9722 \times 10^{24} \text{kg}$ ' or '*Yersinia pestis* is the cause of bubonic plague' or ' $E=mc^2$ '. In debates about scientific realism, it can sometimes be unclear what level of generality we are operating at, or indeed what the correct level at which to operate is. Identifying which level is the correct one in different circumstances will be one project of this thesis.

For now, I will continue to talk as though realism and anti-realism were general approaches to science. This is both because that is how they often have been discussed in the literature and because it will simplify things for the purposes of explaining the two positions.

Section 1.1: The success to truth inference

The most common and widely discussed motivation for scientific realism is the success to truth inference, or no miracles argument (Vickers 2019, Chakravartty 2017, Magnus and Callender 2004). This argument starts by noting that science has been very successful in allowing us to predict, explain, manipulate and intervene in the world. Scientific developments have facilitated accurate and surprising predictions, enabled the development of useful technologies, and so on. The scientific realist reasons that this would all be unlikely if our best current scientific theories were not at least approximately true. Why would an entirely false theory facilitate accurate predictions or useful technologies? However, if our theories are more or less true, it makes sense that they should lead to accurate predictions, or provide the basis for satisfying and wide ranging explanations of the world around us. For this reason, this argument is sometimes called the 'no miracles argument' because, in the words of Putnam, "scientific realism is the only philosophy that does not make the success of science a miracle" (Putnam 1975). The inference grounded by this line of reasoning is the success to truth inference (hereafter STI):

STI: if a theory is very successful, then it is probably at least approximately true

Much of the debate over scientific realism has ultimately come down to the defence or attack of this inference. Anti-realists present historical evidence that this inference has proven unreliable, and the realist spends attempts to save the inference from their attacks. Over time this has led to a number of qualifications and revisions to the basic idea present above, as I will explain below.

Section 1.2: Anti-realism and the Pessimistic Induction

The realist argues that the success of science would be unlikely if much of science was entirely false. Against this, the anti-realist points out that there are many apparent examples in the history of science of theories that were once judged to be very successful, but which are now regarded as entirely false. Laudan (1981) provides an extensive list of such theories, with further examples targeted towards more contemporary versions of realism supplied by Vickers (2013).

Cases such as these appear to undermine the STI. If the STI were a legitimate inference then there should be few cases at most where a theory showed an appropriate degree of success but was not even close to being true. Or rather, a small proportion of all successful theories ought to have been false. However, the historical record, as presented by anti-realists, appears to tell a different story.

Success does not seem to have been any kind of a reliable guide to truth in the past, so we have no reason to suppose it a reliable guide to truth now. This argument is known as the pessimistic induction over the history of science.

Realists have generally taken the PI to be a valid argument, but sought to dispute its empirical premises, i.e., for each putative example of a successful-but-false theory in the history of science they have sought to show it is not as it appears. There are two main ways this can be done. First, one can try to show that an apparently successful-but-false theory from the history of science was not really false, or at least not entirely so. Second, one can try to show that an apparently successful-but-false theory from the history of science was not really successful, at least not in the right way and to a sufficient degree.

Section 1.3: Apparently successful-but-false theories were not really false

There are two ways in which realists have tried to claim that successful-but-false historical theories were not really entirely false. First, realists have tempered the success to truth inference such that it only seeks to infer *approximate* truth, rather than absolute truth (Chakravartty 2017, Vickers 2019). Second, realists have attempted to distinguish the ‘working posits’ of a theory from its ‘idle parts’, and argue that we should only infer the probably approximate truth of a successful theory’s working posits (Kitcher 1993, Psillos 1999).

With regard to approximate truth, the idea here is that although some degree of revision to our theories may in the future be necessary, these alterations will be minor and will leave the gist of our current theories intact. Similarly, realists argue that many apparently successful-but-false theories from the history of science were really approximately true.

Specifying exactly what is meant by approximate is not easy to do in an unproblematic way (Chakravartty 2017). This is one of the reasons why anti-realists tend to be suspicious of the concept. Stanford (2003) argues that establishing that some problematic past theories were ‘approximate true’ after all does not get the realist what they want, i.e., a license to trust the accounts of the world given by current theories. Stanford suggests translating our attitude to current science into a bet about whether the future development of science will resemble its history (Stanford 2015). Stanford contrasts ‘catastrophists’ with ‘uniformitarians’ in reference to the positions taken by nineteenth century geologists. Uniformitarians hold that the future development of science will feature similar upheavals to the transition between Ptolemaic and Copernican astronomy or the transition from Newtonian to relativistic physics. Catastrophists by contrast think that such

upheavals are only a feature of the early development of science and not likely to occur in the future.

There is certainly something right about Stanford's distaste for a realism debate that depends too heavily on how far terms like 'approximately true' should be stretched. For this reason I will not be depending on the idea of approximate truth very heavily in my discussion of the realism debate in this thesis.

The second way in which realists have attempted to argue that historical successful-but-false theories were not really false is through the strategy of selective realism. Here the realist argues that scientific theories can be divided into their 'working posits' and their 'idle parts' (Kitcher 1993). Working posits are those parts of a theory directly involved in generating successful predictions, or other forms of success. Idle wheels are the parts of a theory not involved in generating the theory's successes. They go on to argue that in historical theories that we now regard as false but successful, the theories were in fact partially true. Specifically, those parts of the theories directly involved in generating the important successes were true, while the rest was false. Similarly, selective realists argue that whatever parts of our current theories are needed to generate the successes our current theories enjoy, those parts are at least approximately true.

This line of defence can be attacked on two grounds. First, one might look for examples of cases where it is hard to deny that false parts of historical theories were directly involved in generating success (Saatsi and Vickers 2011 and Vickers 2012 discuss such cases). Second, one could argue that there is no way to identify at the time a theory is our best option which parts of the theory are working posits and which are idle wheels. It is only with hindsight, and with the desired outcome specified in advance, that realists are able to look at historical theories and come up with a story about how the tenets of those theories we still accept were responsible for success. If asked, the realist is not able to say which parts of, for example, the theory of general relativity might be idle wheels. If we did know which parts of our current theories were idle wheels we would probably immediately cease to include them in our theories (Stanford 2003, Vickers 2013). Again, selective realism will not feature heavily in this thesis. It has however been a sufficiently large part of the debate to be worth noting in passing.

Section 1.4: Apparently successful-but-false theories were not really successful

I mentioned earlier that the second way the realist can deal with cases from the PI is by specifying further the kind and degree of success that is needed to infer truth (or rather, probably approximate truth of working posits).

An early move in this direction was to argue that the success to truth inference should only be applied in 'mature sciences'. The idea here was to rule out some of the examples on Laudan's list that appear to have come from a time before anything resembling modern scientific method had been developed, such as the humoral theory of medicine for example.

However, even if there is something in this idea there are a number of much later examples of apparently successful-but-false theories that cannot plausibly be seen as from immature science. Newtonian physics is probably the most obvious such case. Clearly the realist will need further ways to specify what constitutes appropriate success for realist commitment.

Section 1.4.1: Novel Predictive Success

One very important move in this direction has come from qualifying 'success' in the STI to mean 'novel predictive success' (NPS). Realists have argued that there are no cases, or at least very few cases, of theories in the history of science that showed novel predictive success but were not at least partially or approximately true (e.g. Musgrave 1985). The result of this observation for realists has been to move away from defending STI, as given above, and instead attempt to defend the following inference:

Predictive STI (PSTI): if a theory shows a sufficient degree of novel predictive success, then it is probably at least approximately true in its working posits.

What is meant by novel predictive success here? Well, a theory enjoys predictive success if predictions have been made on the basis of the theory that turned out to be accurate. The degree of predictive success provided by a successful prediction is dependent on a few things. Put as simply as possible, these are as follows: First, how risky was the prediction? Second, how novel was the prediction, and in what way was it novel?

How risky a prediction is depends first on how specific it is, and second on how surprising it would be for the prediction to be true if we assume the theory that generated it were false. The first point here, about how specific a prediction was, is fairly intuitive. If someone predicts that it will rain in Scotland in the year 2024, this prediction turning out to be accurate will not be very impressive. If, however, they predict that 6mm of rain will fall in Abernethy between 13:33pm and 14:56pm on 16th October 2024, this prediction turning out to be accurate will be more impressive. A theory that made this second prediction and was vindicated would gain a lot more confirmation than a theory that made the first prediction. Generally, the more specific a prediction, the more its accuracy contributes to the predictive success of the theory used to make the prediction, *ceteris paribus*.

How surprising a prediction is depends on how likely we estimated that prediction to be in the absence of the theory that predicted it. For example, say we develop a new meteorological theory and it predicts that it will rain in Scotland in 2024. That event is near certain regardless of whether we adopt the theory or not, so the prediction offers little confirmation to the theory even when it turns out to be accurate. By contrast, if the theory predicts that the temperature in Scotland will never dip below 20 degrees Celsius in the whole of 2024 and this turns out to be accurate this will provide a good deal of confirmation for the theory. Without the theory we would estimate the chances of this happening as near zero.

Vickers (2019) explains the importance of a prediction's riskiness in Bayesian terms. According to Bayes Theorem, the probability of a theory T given some evidence E is equal to the probability of the E given T, times by the prior probability of T, and divided by the prior probability of E. The more risky a prediction is, the more it increases the probability of a theory being true should the prediction be verified.

The final factor that determines how much confirmation is derived from predictive success is novelty. Novelty is what distinguishes predictive success from explanatory success. Say a theory T1 has an event A as a consequence. Now imagine two different scenarios. First, T1 is developed in advance of A happening and is not designed to account for A. Second, T1 is developed after A happens and is specifically intended to account for A. Intuitively, many people think A provides more confirmation for T1 in the first scenario than in the second. The difference between them is that A is novel for T1 in scenario 1 but not in scenario 2.

There is a whole debate in epistemology about whether and to what extent predictions provide better evidence than accommodations (evidence a theory was in some way built around) (Barnes 2022). The answer to this questions rests to a good extent on what sense of 'novelty' we are working with. Further, some philosophers hold that predictions are inherently superior evidence, while others hold that their superiority is due to their status as predictions serving as a marker of other epistemic virtues. Regardless, there is fairly widespread agreement that, for direct or indirect reasons, predictions are generally better evidence for a theory than accommodations (ibid). Regardless, the motivation for scientific realists to move from STI to PSTI has more to do with trying to get around the PI than it has to do with any abstract epistemological considerations in favour of predictive evidence, although it is important to show that such motivations exist in order for the realist's move not to seem ad hoc (see chapter 2 of this thesis).

What is being done when the realist moves from STI to PSTI is often vague between two options. On the one hand realists may be saying that they think the original STI *is* indefensible, but that PSTI is

defensible. If so this means explicitly denying that non-predictive success can provide adequate grounds for realist commitment. On the other hand, they may simply be refusing to commit one way or the other as far as STI is concerned, but saying that PSTI is defensible. Either way, the result is failing to continue defending STI, and defending PSTI instead. That means failing to defend the idea that non-predictive success (e.g. explanatory success) can provide adequate grounds for realism – it is hard not to feel like this amounts to a tacit concession to the anti-realist that non-predictive success cannot be a guide to truth (even the approximate of working posits). Whether that is the intention or not, the almost exclusive focus on PSTI among realists leaves the territory of non-predictive success as a guide to truth occupied only by anti-realists arguing that STI does not work. As I will be arguing later, this is not a satisfactory situation.

Returning now to the defensibility of PSTI, this version of the inference still faces major historical challenges. Vickers (2013) mentions a number of cases of false theories generating NPS. However, it still seems that focusing on NPS rather than a more general notion of success is one important tool in the realist's arsenal. Modern statements of the STI still tend to be phrased only in terms of NPS (e.g. Vickers 2019).

A problem for this move to focusing on predictive success alone is that there are areas of science where novel predictive success appears not to play a major role. By focusing only on PSTI the realist appears to abandon realism about those areas. I will return to this problem, and develop it in significant detail in chapters 2 and 3.

Summary of Thesis Argument by Chapter

The overall argument of this thesis is that the historically focused scientific realism debate described above goes wrong by conflating many different methods and types of evidence under the general heading of 'success' and also fails to recognise that the reliability of any method or type of evidence varies by context. I argue that we should think of historical inductions on the history of science as empirical tests of aspects of scientific method. Such historical inductions only represent legitimate tools for evaluating scientific methods if they really do concern one method, and further if they only try to judge that method for relevantly similar contexts.

In chapters 2 and 3 I begin by arguing that novel predictive success based versions of the success to truth inference are arbitrarily narrow and poorly motivated. The role of these chapters is partly to bring out a shortcoming in the existing scientific realism debate (and therefore motivate the idea that a different approach to the debate is needed) and partly to illustrate the diversity of scientific

evidence. This will begin to undermine the idea that all possible grounds for realism about a theory can be captured by any version of the success to truth inference.

The argument of chapter 2 is that the concepts of prediction and explanation are too closely related for it to be plausible that the legitimacy of a predictive success to truth inference really is independent of the legitimacy of a non-predictive success to truth inference. To maintain that the two inferences really are independent one would need to hold to the extremely implausible view that accommodated evidence provides no confirmation for a theory whatsoever, while predictions provide at least some.

In chapter 3 I develop further the idea that the view of non-predictive evidence needed to maintain the independence of the predictive success to truth inference is unsustainable. I do this by looking at the history of evolutionary biology from 1859 to 1869. I first look at how it came to be accepted by scientists that evolution was a fact of the natural history timeline – that is, I look at how it came to be accepted that current species are the result of descent with modification from earlier species. I then go on to look at how it came to be accepted that natural selection is the most important factor in explaining evolutionary change. I argue that in both cases the degree to which novel predictive success was a crucial consideration in assessing the strength of these theories was limited. Much of the important evidence cannot plausibly be regarded as novel predictive success. For much of the evidence that can be understood as novel predictive success, either its status as predictive has little to no epistemic importance or the predictive evidence is insufficient to explain the strength of the overall case for these theories. This should teach us that novel predictive success cannot be the whole story when it comes to what should decide our realist commitments, and also that the realist's attempt to defend PSTI in isolation from the more general STI are poorly motivated.

In chapter 4 I begin to work towards an alternative approach to the historically focused scientific realism debate. I argue that historical inductions on the history of science are best understood as exercises in methodological naturalism – i.e., the idea that our views of scientific method should be justified by empirical testing. I outline some potential problems for this approach and argue that they can be overcome, and argue that seeing things this way maintains the heart of the traditional historically focused scientific realism debate. I also begin to explore some issues that this way of seeing things raises. In particular I argue that if we see the success to truth inference as a methodological principle then it becomes reasonable to ask how well it describes the actual methods of science. I argue that the plausibility of any global approach to the realism debate rests on the degree to which universalism about scientific method can be justified. I further argue that historical assessments of methods generally only work when they are about the reliability of one

method for one topic. Again, this ties any global approach to the realism debate into thinking of the success to truth inference as a reasonably unitary method and thinking of 'science' as a reasonably unitary topic.

In chapter 5 I argue that universalism about scientific method is not viable. I note the general unpopularity of universalism and argue against it as a view of scientific method in quite general terms. I then go on to ask what version of a historically focused realism debate might be compatible with the insights of methodological pluralism. I distinguish two sources of plurality in scientific method: one based on the existence of multiple sets of methodological rules, and one based on the context sensitivity of scientific methods. I present some serious difficulties with reconciling either of these forms of pluralism with historical assessments of scientific methods. Versions of methodological pluralism inspired by the context sensitivity of methods are particularly hard to reconcile with a historical realism debate.

In chapter 6 I outline and discuss recent work on localism in the scientific realism debate. Localism is any version of the view that whether we should have realist commitment to scientific theories should be decided 'locally', for example at the level of individual theories, or perhaps at the level of different disciplines. The view I am working towards is a version of localism about the realism debate, since I advocate engaging with historical inductions about the reliability of methods at a level far below the 'global' level of science as a whole. I present arguments made by Asay (2017) and Fitzpatrick (2013) in favour of different sorts of localism, together with criticisms from Park (2019) and Henderson (2018). I use these philosophers' work to present a tension between the way localist motivations seem to pull us towards the view that each theory should be evaluated individually, and the need to compare contentious pieces of scientific reasoning to a generally approved form of reasoning in order to justify them. I argue that we can access these comparative justificatory resources so long as we do not think that scientific methods are so context sensitive that each theory must be evaluated independently. I point out that holding that scientific methods are highly context sensitive does not necessarily avoid historically motivated scepticism, as the anti-realist can argue that we have proven to be poor judges of whether a method applies in a given context. I argue that in order to avoid this kind of argument, we need to be able to categorise contexts into types. If we can identify what type the context we are in belongs to, and can show that a method has proven reliable in that type of context, then this gives us grounds to reject the idea that we are poor judges of whether the method in question applies in the context we are in.

In chapter 7 I look at how we might categorise types of scientific context, and argue that we should categorise scientific contexts according to the types of difficulty present in them. The barriers to

knowledge vary across scientific fields – for example, some fields are difficult because they concern things too small to observe, while others are difficult because they concern highly complex systems. I point out that anti-realists' attempts to delineate the limits of their scepticism always seem to be versions of the idea that we should be sceptics about some parts of science because those parts of science are too difficult. However, anti-realists characterisations of what this difficulty is and which parts of science it is found in are often lacking some refinement. I attempt a catalogue of different types of difficulty, and argue that the reliability of any method for each type of difficulty is largely independent of its reliability for other types of difficulty. The view of the scientific realism debate I end up with is one where the history of science can be used to test the track record of a method relative to a type of difficulty. I explore the implications of this for the realism debate. I argue that since in any actual scientific context there will often be multiple methods and types of difficulty in play, and what these are will not be obvious without detailed local knowledge of a theory, it may often be true that whether to be a realist about any particular scientific theory should be decided on an individual basis for each theory.

In chapter 8 I conclude my thesis by recapping the most important points of my position, and comparing my position with some other writers on localism in the realism debate. This comparison is meant to bring out how my work adds to the range of options available in this area. I argue that the most significant aspect of my thesis is that, if successful, it reconciles the insights of those who point to the context sensitivity of scientific confirmation and the need to engage with each theory on its own merits, with the possibility of learning substantial epistemic lessons from the history of science that can be relevant to deciding which parts of science we should be epistemically committed to.

Chapter 2: Can the predictive requirement on the Success to Truth Inference be independently Motivated?

Section 1: Why independent motivation is needed

A requirement in the scientific realism debate is that when the realist suggests revisions to the success to truth inference, these revisions should have independent motivation and not be an ad hoc attempt to save the inference by whatever means necessary. This means it is not enough just to suggest a revision to the success to truth inference because it gets around some problematic historical cases. For example, say there were no examples in the history of science of theories that were successful-but-false which were invented by someone named Smith. This would not allow us to create an inference of the form 'if a theory is successful and was invented by someone named Smith then it is probably approximately true'. This inference might be immune to the pessimistic induction, but there is no reason to think the name of a theory's inventor is relevant to its epistemic status. This addition to the inference lacks any independent motivation besides its convenience in getting around the pessimistic induction.

When the realist qualifies 'success' in the success to truth inference to mean 'predictive success', this move requires independent motivation. We need some general epistemological argument to the effect that this move is not arbitrary or ad hoc. Otherwise, even if it should happen that there are no predictively successful but false theories in the history of science, there is little reason to think this saves realism from the pessimistic induction. The anti-realist can respond to the realist's claim that there are no cases of predictively successful-but-false theories in the history of science by arguing that there is no reason to think this more than an accident.

Section 2: What would it require for us to say there was independent motivation for specifying that 'success' means predictive success?

In order for the NPS qualification on the STI to have independent motivation the realist needs to show two things. First, the realist needs an argument to show that the validity of the NPS based STI is independent of the validity of a version of the STI that includes non-predictive success. Second, there needs to be independent reason to think that predictive success is generally epistemically superior to other kinds of success. I expand on each of these requirements in turn below.

Independence: Consider the following three versions of the success to truth inference.

1. Success to truth inference (STI): if a theory shows sufficient success then it is probably approximately true.

2. Predictive success to truth inference (PSTI): if a theory shows sufficient novel predictive success then it is probably approximately true.
3. Non-predictive success to truth inference (NPSTI): if a theory shows sufficient explanatory scope and power, coherence with other areas of knowledge, fruitfulness for further research and so on, then it is probably approximately true.

The difference between them is that where STI is happy to infer truth from any kind of success, so long as it is present to a sufficient degree, PSTI only tries to infer approximate truth from the fact that a theory has successful predictions, while NPSTI only tries to infer approximate truth from the presence of other forms of success.

In the previous chapter we saw how realists have tried to overcome the pessimistic induction by moving from STI to PSTI. However, for this move to add anything to the debate it must be the case that PSTI and NPSTI are independent from one another. What I mean by that is that it must be possible that one is legitimate at the same time that the other is not. Imagine that the two are not independent – imagine that if PSTI is legitimate then NPSTI must also be and vice versa. In that situation the existence of any kind of successful-but-false theories in the history of science undermines PSTI even if there are no examples of cases that are specifically predictively successful-but-false. Were that the case, no defence against the pessimistic induction would be achieved by moving from STI to PSTI.

The epistemic superiority of predictive success: the second thing we need for the move from STI to PSTI to be independently motivated is some kind of argument to the effect that successful predictions provide better evidence for a theory than, for example, explanatory success. After all, moving from STI to PSTI is supposed to strengthen our position against the pessimistic induction. We are supposed to now be better protected against inferring (probable approximate) truth on the basis of inadequate evidence than we were previously. For us to be in that position after moving to PSTI, it must be the case that predictions are better evidence than other forms of success.

So long as PSTI and NPSTI are independent, just showing some degree of superiority for predictions is adequate. However, as we will see later, if they are not independent then a much higher, and harder to defend, degree of superiority for predictions is needed.

Is the predictive success inference independent from the non-predictive success inference?

The first task I set the realist was that of showing that the legitimacy of the predictive version of the STI is independent of the legitimacy of the unconstrained STI. I will be arguing here that this cannot be done.

If the two questions are independent, then all of the following combinations of answers should be possible options.

	PSTI	NPSTI
Option 1	Valid	Valid
Option 2	Not valid	Not valid
Option 3	Not valid	Valid
Option 4	Valid	Not valid

Options 1 and 2 are clearly possible, but if *only* they are then the questions are not independent. The real test for whether the two success to truth inferences are independent is whether option 4 is at all an option (I don't think anybody is likely to be tempted by option 3, so I will ignore it). If the validity of the two inferences are independent of one another then it should be possible to simultaneously hold different evaluations of each one.

It is worth expanding on what I mean by 'possible' here. Clearly options 1-4 cannot all be right; at most one of them can be. By 'possible' I mean that we cannot know in advance of looking at the history of science which one is true. What should ideally decide the choice between options 1-4 should be surveying the history of science and seeing first whether PSTI can withstand the pessimistic induction, and second whether NPSTI can withstand it, and then seeing which combination of answers we have ended up with. However, if there are reasons to think that an option is in some way inherently unstable or unsatisfactory whatever the history of science may show us, then that option is not a 'possible' contender in the way I am using the term.

I do not think option 4 is possible in this sense, and for that reason I do not think the two inferences are independent. First, I will argue that predictive and non-predictive success are too closely related for it to be very plausible that their evidential merit could be independent to the extent that option 4 requires. Second, and more crucially, I will argue that option 4 has very implausible consequences in scenarios where we contrast the degree of confirmation between one theory with a lot of non-predictive success but no predictive success, and another with little non-predictive success and a very small amount of predictive success.

Section 3: Why prediction and explanation are related concepts

The main form of non-predictive success is explanatory success. A theory enjoys explanatory success if it is able to explain a wide range of phenomena in a simple yet powerful way. What does it mean for one thing to explain another in science? This is a question that has attracted significant debate,

but something all accounts of scientific explanation have in common is that they take part of what makes it true that A explains B is that B is what we would expect given A. This idea of something being explained if it is what we would expect given our current knowledge seems central to the concept of explanation. However, to say that B is what we would expect is the same as saying that B is what we would predict. There is a fundamental link between prediction and explanation in that whether a theory A predicts B and whether it explains B depends on the same thing – whether it is a consequence of A that B.

For this reason, the very same piece of evidence can be either something that is explained by a theory or something that is predicted by it depending on when that evidence happened to come to light, and what concept of novelty we are working with when deciding whether the evidence is a prediction or not.

A final point to make about the link between prediction and explanation is that in almost any actual situation where we are comparing theories and trying to choose between them, predictive and explanatory successes for each theory will be balanced and compared side by side. Both will fall generally under the category of 'evidence' as we tot up the pros and cons of each theory.

All of this means it would be very surprising if predictive and non-predictive (explanatory) success were to diverge quite *as* wildly in their epistemic import as option 4 would have us think.

So far of course though, this is not a conclusive argument against option 4 as a possibility – just reason to think that its prior probability is low. In the next section I develop my argument against the possibility of option 4 further.

Section 4: Theory choice when we have to balance predictive and non-predictive success

Imagine we have two theories, T1 and T2. T1 has a lot of very impressive explanatory success, coherence with other areas of knowledge, fruitfulness for further research etc. but no predictive success of any great importance. T2 has much less explanatory success (maybe none even) and one impressive predictive success. Option 4 would lead us to say that T2 is better confirmed than T1. Or at least that it is a better candidate for realist commitment, which amounts to the same thing. On option 4 we can only infer (probable approximate) truth from predictive success, so we have to say that T2 is better confirmed than T1, no matter how good the non-predictive evidence for T1, and how unconvincing the general case for T2.

The obvious defence for the option 4 advocate here would be to say that despite its lack of intuitiveness there is nothing necessarily wrong with the situation described above. If we have to, we can just say that T2 really would be better than T1. This argument could perhaps be enhanced by

denying that in actual situations resembling the one I describe our intuitions really would favour T1. Many of the most readily available examples of theories like T1 are to be found on the lists of the pessimistic induction, so perhaps it is just a surprising lesson of the pessimistic induction that any amount of predictive success trumps any amount of non-predictive success.

The first part of my reply to this defence is to say that the realist is now in a very difficult position with regard to independently motivating the idea that prediction is epistemically superior to non-predictive evidence. Just showing that prediction is generally better evidence than, for example, explanatory success is probably achievable. However, that is no longer enough. To make option 4 work we need independent motivation for the much stronger claim that predictive evidence always, under any circumstances, trumps non-predictive evidence. There is little reason to think this will be achievable.

The second part of my reply to option 4's defence will be to show that in at least some important cases there have been very well confirmed theories in which predictive evidence has not been important. The cases I will be discussing will be the debate over whether modern species have evolved or were created, and the debate over what mechanisms lie behind evolutionary change. I aim to show that predictive success was not an important part of the case either for evolution as a fact, or natural selection as its explanation. To say that the theories of evolution or natural selection are less well confirmed than a theory with only a small amount of predictive success and no non-predictive success to its name is not plausible.

Section 5: Independent Grounds for thinking predictive success is epistemically superior

The task of showing that we have independent grounds for thinking predictive evidence superior to non-predictive evidence is probably achievable. The view that predictive evidence is stronger than non-predictive evidence is called 'predictivism'. Most of the work done on this issue contrasts prediction with 'accommodation' (Barnes 2022). An accommodation is a fact that a theory is somehow built around, whereas a prediction is a consequence of a theory that the theory was not built around. A prediction makes a claim that is in some way new, or at least new relative to the theory in question. Hence the 'novel' in 'novel predictive success'. There are a variety of ways to characterise this novelty. Temporal novelty is perhaps the most straightforward. A prediction is temporally novel if it was not known to be true at the time a theory was constructed. Other accounts of novelty rest instead on whether a fact was used in the construction of a theory or not, and whether a fact was a consequence of any rival theory at the time the theory was invented (ibid).

A number of philosophers have argued for some version of predictivism, and most seem to have settled on some form of weak predictivism – the view that predictions are superior to accommodations not inherently but because they indicate the presence of other epistemic virtues. If PSTI and NPSTI were independent, this would be enough. However, to hold that predictive evidence can be a guide to truth and non-predictive evidence cannot (i.e. option 4 above) requires something stronger. It requires us to show that accommodated evidence provides *no* confirmation for a theory, but that predictive evidence provides some. That view is known as the null support thesis. Despite a brief period of popularity in the 1980s the view no longer has serious defenders (ibid).

The null support thesis was advocated separately by Giere (1984), Glymour (1980) and Zahar (1983) on the grounds that accommodations are facts that a theory was in some way designed to fit. Since the theory was designed to fit those particular facts, there was never any chance that it would not fit those facts – there was no risk of it being falsified by them. The theory therefore gains no evidential support from accommodated facts.

The flaw in this argument is that either some piece of evidence is a consequence of a theory or it is not, regardless of whether the theory was designed to fit that evidence or not. Even if a theory was designed after some experimental result became known, there was still a risk that the experimental set up in question could have falsified that theory, as it could have happened that the experimental set up gave a different result. Had it done so, it would have refuted the theory. This argument was given by Howson (1990).

To illustrate further the problem with the null support hypothesis, in some circumstances it seems like a legitimate method to go and look for some facts and then think what theory suggests itself on the basis of those facts. For example, a detective might look at a crime scene and find that the butler's fingerprints are on the murder weapon. He might then theorise that the butler did it. Although this theory is designed to fit facts that are already known, it makes very little sense to think those facts provide no confirmation for the theory. There is nothing about the act of looking for fingerprints on the murder weapon that fails to offer a chance of falsification to the theory that the butler did it. There was every chance that we could have in fact found the victim's heir's fingerprints. When he talks about the evidence not having a chance to falsify the theory, Giere confuses the experimental setup with the result.

The null support hypothesis is rightly unpopular. The fact that option 4 requires it to be true is a big mark against option 4. There is nothing good to say by way of independent motivation for the position taken by option 4.

Section 6: How much does the move to a predictive success to truth inference really need independent motivation?

It looks like the only possible route to the idea that predictive success can ground realist commitment even if non-predictive success cannot is through the pessimistic induction. Grant for the sake of argument that the pessimistic induction shows that non-predictive success is not a reliable guide to truth, but is silent on whether predictive success is a reliable guide to truth. Could we see this as providing empirical evidence for option 4? It might be that we cannot independently motivate the null support thesis through abstract epistemology, but perhaps we can learn it as a lesson of history? Maybe the demand for independent motivation for revisions to the success to truth inference was misconceived. Or at least maybe we don't need independent motivation for this particular revision.

Of course, there is nothing in that line of thought to convince an anti-realist. It seems like an at least equally appropriate response to the position we find ourselves in here to say that the existence of successful-but-false theories in the history of science undermines both the original version of the success to truth inference and the predictive version of it, since no independent argument can be given that the two inferences are independent. The anti-realist is still justified in their suspicion that the realist's qualification on the inference that only predictive successes are to be counted is ad hoc.

If there were no other argument against the viability of option 4 then the debate would perhaps be a stalemate at this point. However, in the next chapter I will use case studies to develop a positive argument that there is something wrong with preferring T2 in the scenario I sketched earlier. There are some very well established T1s in the history of science. The problem is not just that the split between the predictive and non-predictive versions of the STI cannot be independently motivated – there is positive reason to think the split is illegitimate.

Chapter 3: Case studies in non-predictive success: history of evolutionary biology 1859-1937

Introduction

In the previous chapter I argued that no independent support can be given to the realist's move from STI to PSTI as it is difficult to defend the claim that PSTI and NPSTI are independent. Doing so requires committing to some implausible views about how theories compare when one enjoys a very small amount of predictive success and another has no predictive success but a very large amount non-predictive success. Those implausible views in turn rely on an epistemic principle that is very hard to find independent motivation for, namely the null support thesis. The result is that the move from STI to PSTI cannot be independently motivated. That is a problem as the availability of independent motivation besides the desire to avoid the pessimistic induction is generally taken to be a requirement for amendments to the success to truth inference.

All of this left the move from STI to PSTI in a tricky position. However, I noted that the realist could perhaps cling to the idea that there is no need to motivate the move from STI to PSTI in a way that is *independent* of the debate over the pessimistic induction. Perhaps if that is what the historical record tells us we could just take it to be a surprising lesson of the pessimistic induction that predictive evidence (however small) always trumps non-predictive evidence (however significant). Even if we cannot get there through abstract epistemology, maybe we can through the history of science.

In this chapter I will argue that this idea is no good. There are some very well confirmed theories the case for which is not crucially concerned with novel predictive success. Realism about these theories is as plausible as realism about any theories in science. It just does not make sense to say that a theory with only a very small amount of predictive evidence in its favour is better confirmed than the theories I will discuss.

The theories I will look at will be the broad theory of evolution, and the theory of natural selection as one of the primary mechanisms of evolutionary change. In the case of the theory of evolution, I will argue that although some of the evidence in favour of evolution could be seen as novel predictive success, there is no interpretation of 'novel' on which this evidence was both novel predictive success, and on which its status as such can be given any epistemic significance. In the case of natural selection, I will argue that the key scientific developments that allowed for the widespread acceptance of natural selection as the primary mechanism driving evolution should not be thought of as NPS for the theory of natural selection. Where the evidence for natural selection can be characterised as NPS, I will argue that either this evidence did not play a crucial role in the

overall case for natural selection, or its status as a prediction is unimportant for explaining its epistemic significance.

The overall lesson I take from these episodes in the history of biology is that NPS cannot be the whole story for the realist. Building on the work of the previous chapter, these episodes show us that we should not see NPS based versions of the STI as independent of more general versions of the STI and should not rely on the concept of NPS too heavily in any defence against the historical anti-realist. The history of evolutionary biology in this period also serves as an illustration of the importance of scientific context to determining the strength of any kind of evidence, and the complexity in the case for any real scientific position. There are almost always likely to be different sorts of evidence in play, where these cannot easily be grouped under a single category such as NPS or even 'success' more generally. These points about diversity in scientific contexts and sorts of evidence will be discussed further in the remainder of this thesis.

Section 1: Historical introduction

After the publication of Charles Darwin's *Origin of Species* in 1859, biologists came to accept fairly quickly that modern species had evolved from very different earlier forms through descent with modification. Within about 10 years most seem to have been convinced, and certainly by 1900 very few scientists still defended special creation as the explanation of the origin of species (Bowler 1985). In part 1 of this chapter I will outline the reasons why evolution was accepted, and analyse what lessons can be learnt from this about the relative importance of predictive success and other kinds of evidence. I will argue that in this case the distinction is either unclear or unimportant.

Unlike with the fact of evolution, scientists were less quick to converge on Darwin's proposed mechanism of evolutionary change, i.e. natural selection. This did not happen until the modern synthesis between Darwinian natural selection and Mendelian genetics came about, which cannot certainly be said to have happened until the publication of works like Dobzhansky's (1937) *Genetics and the Origin of Species* (ibid). In the intervening time a variety of other explanations of evolutionary change were more popular at various points – particularly the position known as 'mutationism'. I will outline the reasons first why natural selection was controversial, and second why it eventually triumphed, before analysing what lessons this episode in the history of evolutionary biology gives about the distinction between predictive and explanatory success and their relative importance.

Section 2: Evolution

There were three main reasons why evolution was quickly adopted after 1859. First, it was widely regarded as the only scientific option. Second, it had significant advantages over creationism in terms of its ability to explain various phenomena in biology. Third, important intermediary fossils, as well as sequences of invertebrate fossils, were found which seemed to show evolution in action.

Section 2.1: Evolution as the only scientific option

When the *Origin of Species* was published in 1859 the view that species were each created in their current forms by divine intervention was still widespread. One drawback of this was that there wasn't really anywhere for a scientist to go. There is no point in trying to *explain* the forms, distribution or behaviour of animals if the answer to any question that might be asked really comes down to 'because God chose to make it so'. Cataloguing the various species that exist was possible, as was *describing* their anatomies, behaviours, distributions and relationships to one another. But a framework in which such things could be explained was lacking. An early convert to evolution, Asa Gray, made this point when he argued that in terms of biology creationism 'neither explains nor pretends to explain' (Larson 2004 p.88).

Bowler (1985, p.177) and Larson (2004, p.88) both agree that the scientific nature of evolution as opposed to special creation was a major factor in biologists adopting it. Larson writes: "Everywhere evolutionism took root, it held a similar appeal for scientists. With a theory of evolution, laboratory biologists and field naturalists could begin trying to explain the origins of living things (and perhaps of life itself) in terms of regular, rational, repeatable natural processes rather than divine fiat." (ibid)

Thinking of the origins of species in terms of evolution opened up scientific questions and research projects that had not previously existed, as well as giving others that had previously existed new importance. This fruitfulness of the theory of evolution was an important part of why scientists were keen to adopt it. Larson cites Alfred Russell Wallace's work on the geographical distribution of species, and Ernst Haeckel's work in embryology as examples of projects inspired by evolutionary thinking. Another obvious example lies in the search for intermediary forms both in the fossil record and in extant species that followed Darwin.

Section 2.2: The explanatory power of evolution

Things that evolution offers a satisfying explanation of, but which special creation is unable to shed light on include:

- Homologies. Homologies are structural similarities between different species. For example, the flippers of whales, the hands of humans, the wings of birds, and the forelimbs of dogs all show remarkable similarity of structure – each share many of the same bones arranged in

much the same way. This had been observed before Darwin, most notably by Owen, who coined the term 'homology', however had no satisfactory explanation. (Bowler 1985 p.182) If all the species with the homologous structure in question evolved from a common ancestor however, this would explain the phenomena. Each has adapted the body plan of the common ancestor to serve the needs of their environment.

- Rudimentary organs. Also well explained by the theory of evolution is the presence of rudimentary organs in many species that do not have use for those organs. For example, snakes and whales have rudimentary hip joints despite not having legs. These make sense if one thinks that snakes and whales evolved from ancestors that did have legs, but makes little sense if one thinks they were created in their current form (ibid).
- Biogeographic distributions of extant species. The distributions of species across the globe was taken as another thing that evolution could explain better than creationism. (ibid p.180) Different regions of the globe have their own distinct collections of flora and fauna. The creationist explanation of this was that God had chosen to create each species in the place suited to it. However, not all of the facts of distribution made sense on this picture. Similar habitats in different areas of the globe can have almost entirely different species, and more widely differing habitats can often share some of the same species as long as they are in the same biogeographic zone. This is better explained by the impossibility of crossing certain physical barriers for large chunks of the development of modern species. A striking example of this is Wallace's line in the Pacific, on one side of which are found Australasian species, and on the other side of which are found Eurasian species. There is no great difference in conditions either side of this line, but the impossibility of travel across it over the course of evolutionary history means that completely different species occupy each ecological niche.

Section 2.3: Intermediary fossils

If species have evolved by descent with modification, one would expect to find fossils of intermediary forms between extant species and known extinct species. Were such fossils to be found, it would lend considerable support to the idea that species had evolved, since there was no reason for a creationist to expect to find a creature that was halfway between a bird and a reptile, for example.

An early reaction to the *Origin of Species* was therefore to begin looking for intermediary fossils. This project had several successes in the nineteenth century.

One important discovery was of *Archaeopteryx*. *Archaeopteryx* was a species that lived around 150 million years ago, had feathers and large wing-like forelimbs like a bird, but had a mouth and teeth

rather than a beak, and many other features in common with other dinosaurs. The first complete *Archaeopteryx* specimen was found in 1861. (ibid p.192)

Discoveries such as these do seem to have played an important role in convincing scientists in the late nineteenth century. Bowler claims that discoveries of sequences of invertebrate fossils showing a transition from one form to another were what convinced most palaeontologists that species evolved. (ibid p.193)

Section 2.4: Discussion

At first glance, the discovery of intermediary fossils looks like a novel predictive success for the theory of evolution, and also looks like it played a crucial role in convincing people that modern species came about through descent with modification.

Intermediary fossils were something that the theory of evolution would lead us to expect, but which creationism would not. Intermediary fossils were found after the most influential statement of the theory of evolution, so clearly were temporally novel as well as every other kind of novel. So in this regard they seem fairly clearly like predictive evidence.

One thing worth considering is how much attention we should pay to the fact that the prediction of intermediary fossils was not at all quantitative. There may be grounds for thinking that the prediction was not 'risky' in the appropriate way. No prediction was made concerning when, where or how the fossils would be found. No prediction was made concerning the exact nature of the intermediary fossils. So the prediction was certainly vague. However, there does still seem to have been a significant degree of risk involved in the prediction because there was no reason to expect it to be true, in spite of its vagueness, if it were the case that species had been specially created. The prior probability if we assume that descent with modification is false is low.

A bigger consideration is whether it would have had any importance epistemically if it had so happened that the impressive sequences of invertebrate fossils, and specimens such as *Archaeopteryx* had been found before descent with modification had been suggested. It is hard to see why it would. The evidential value of fossil sequences seems to have little to do with their status as predictions. The link between evidence and theory is so clear in this case that had the theory of evolution not been suggested before they were found, it is hard to imagine it would not have been suggested quite soon after. When we look at intermediate fossils, and especially when we have a complete sequence of gradually changing fossils, it is hard not to feel we are *seeing* evolution happening.

But then perhaps temporal novelty is not what we should be concerned with. As I mentioned earlier, those who advocate the epistemic superiority of predictions generally characterise the novelty that makes something a prediction in non-temporal terms (Barnes 2022). Do these other ways of understanding what novel predictive success is help us?

One alternative conception of novelty is use novelty (Barnes 2022). Some piece of evidence is use novel for a theory if the evidence was not used in the construction of the theory, or if the theory was not built to fit that evidence. However, what I said about the significance of temporal novelty also seems to apply to use novelty in this case.

Another conception of novelty is one on which evidence is novel for a theory if that evidence is not predicted by any known competitor of the theory. However, in the current case there are only two competitors present – descent with modification and special creation. By this definition any evidence at all that might distinguish these competitors would have to be regarded as novel predictive success, so this no longer seems to capture what was supposed to be important about prediction.

What general lessons can we learn from this case about the relative importance of predictive and explanatory success? It is tempting to say that we have seen that predictive success is not necessarily epistemically superior to explanatory success, however this is probably too quick. It may be that evolution is a special case because there are only two (broadly understood) alternative theories of the origin of species – special creation and descent with modification. Usually, this either would not be the case, or we could not be confident it was the case (because of the possibility of unconceived alternatives, to be discussed in chapters 6 and 7). However, even if that is what is going on, that does at least show that global predictivism is not true, in which case there are at least some situations where the novel predictive success based success to truth inference cannot help us reply to the pessimistic induction. In those situations we need to figure out some other response if we want to maintain realism about any theories in those areas. Since realism about the fact of evolution is as plausible as realism about any scientific theory, figuring out what that other response to the pessimistic induction might be would be desirable.

I will now move on to consider the next chapter in the history of evolutionary biology. First the ‘eclipse of Darwinism’ in the late nineteenth and early twentieth centuries, and second the modern synthesis that eventually vindicated natural selection. In this period biologists debated and eventually settled on an answer to the question of what the main mechanisms of evolutionary change are. In this debate there were more than two competitors so in this respect the situation differed from the debate over whether evolution had happened or not.

Section 3: Natural Selection

Section 3.1: Why was natural selection not widely accepted in the nineteenth century?

There were a variety of reasons why natural selection was not popular in the nineteenth century, and not all of them were scientific. For instance, the lack of purpose or the guarantee of progress offered by natural selection was distasteful to many Victorians. However, I will pass over these issues to focus on the scientific objections to natural selection in the years following the *Origin of Species*.

Three of the main problems facing selectionism were Kelvin's estimates of the age of the earth, the evidence then available from the fossil record, and the lack of understanding of variation and heredity (Bowler 1985, Larson 2004). The last of these was probably the biggest issue.

Age of the Earth: for natural selection to be plausible as the mechanism by which species evolved into their current forms it was necessary that the world be thought of as very old, since evolution by natural selection would be a slow process. In 1868 Kelvin published calculations concerning the rate at which the earth had cooled which gave the conclusion that it could not have been more than 100 million years old (Bowler 1985 p.194). This was not seen as a plausible amount of time for evolution to have occurred through natural selection. This concern motivated scientists to look for other, faster mechanisms by which evolution could have occurred such as the inheritance of acquired characteristics (Lamarckism) or mutationism, which held that evolution was primarily driven by sudden mutations rather than by the gradual working of natural selection.

The fossil record: at the time Darwin published the *Origin* the available fossil evidence appeared to show species remaining unchanged for the duration of their existence before suddenly disappearing and being replaced by quite different species. On a gradualist view of evolution like Darwin's theory of natural selection, this is not what one would expect. Instead one would expect to see gradual changes in the forms of species over time.

Darwin's answer to this issue was to point out that the fossil record is by its nature very incomplete, and was particularly so in 1859 (Darwin 2003 [1859] chapter 9). While true, many at the time did not find this convincing. This led some to hold on to a view of repeated catastrophes followed by bouts of special creation, and others to lean towards accounts of evolutionary change that allowed for sudden leaps in evolution, such as saltationism.

Inheritance and variation: the biggest scientific challenge facing natural selection came from the fact that very little was known about how inheritance works.

To understand this problem, it is worth recounting the structure of the explanation natural selection offers for evolution. The account goes as follows:

1. Different organisms in a species exhibit variation in traits relevant to their ability to produce offspring in a given environment.
2. This variation is heritable. For example, a fast runner is likely to produce offspring that also run fast.
3. If a trait (e.g. fast running) confers an advantage to procreation in a species' environment, that trait will tend to become more widespread in the population. For example if fast runners produce more offspring, and these offspring will inherit their parents' ability to run fast, then the average running speed of the population will increase over time.
4. Over time the accumulation of these kinds of shift in the traits of a population caused by natural selection can lead to significant changes in the phenotype of a species, and eventually can give rise to new species.

If each of these premises is true, the account provides a convincing mechanism for evolution. However lack of knowledge concerning the sources and extent of variation and the mechanisms of inheritance at the time Darwin published led to various concerns with this general picture. At the time Darwin wrote, the extent to which (1) above was true was not known. Whether (2) was true was not known. In the early development of genetics many scientists did not believe (3) to be true. Further, many scientists in the late nineteenth and early twentieth centuries doubted whether (4) could be true or whether there would be a limit to the degree of change that could be caused by natural selection. Every step of the account given above was at some point brought into doubt as scientists searched for a workable theory of inheritance and variation.

In 1868 Fleeming Jenkin published a review of the *Origin of Species* in which he raised concerns about the ability of the effects of selection to accumulate in a way that would lead to significant phenotype changes (Jenkin 1868). Jenkin argued that natural selection could only work on continuous variation in a population. For example, if height conferred a breeding advantage then it may be that natural selection would cause the population to become taller over time. However, this trend would be limited by the initial range of variation. Jenkin believed that selection on the normal range of continuous variation alone would not be sufficient to give rise to new species.

Jenkin went on to argue that if an individual were born with very different traits to the rest of the population, different enough to provide material for selection to act on that might eventually give rise to a new species, they would not be able to make a significant difference to the general makeup of the species, even if they were at a selective advantage. This is because their traits would be

‘swamped’ by them and their offspring breeding with the rest of the population. This argument was particularly forcible since at the time most biologists had a ‘blending’ view of inheritance in which the characteristics of both parents would be blended in the offspring.

Darwin and Wallace both recognised the strength of Jenkin’s concern over the swamping of traits possessed only by a very abnormal individual, and saw that selection did indeed need to work with a continuous range of variation across a population (Bowler 1985 p.199).

Worries about the limits of what could be achieved by selection acting on a continuous range of variation in a population continued. Galton believed that any trait would eventually regress to the species mean, even if lying at either end of the range did confer a competitive advantage (Bowler 1985 p.240). And the most convincing case for the idea that there would be limits on what natural selection could achieve was provided by Johannsen’s experiments on beans (Johannsen 1955 [1909]). Johannsen carried out breeding experiments on the common bean, *Phaseolus vulgaris*, in which he self-fertilised an initial 19 plants and bred their descendants over several generations. In each generation, Johannsen operated selection based on the size of beans. He chose either large or small beans over several generations and looked to see if selection had an effect on the mean size of beans produced by subsequent generations. He found that it did not. Within each ‘pure line’, i.e. plants all descended by self-fertilization from one plant, both the range of variation and the mean bean size was not significantly altered by selection across generations.

Johannsen took this as evidence that all plants in each ‘pure line’ belonged to the same ‘type’. Johannsen would eventually develop the concept of ‘genotype’ as being the thing that the plants in each pure line shared. An organism’s genotype is its genetic makeup and its phenotype is its physical characteristics. Johannsen found that a continuous range of variation in bean size was found in each generation purely as a result of environmental factors. However which ‘type’ each plant belonged to was discontinuous – they either belonged to one type or another, and this was what served best to predict the average size of the beans they produced. Johannsen concluded that natural selection can sort between different ‘types’ in a genetically diverse population, but cannot give rise to new types. This gave convincing evidence to scientists who believed that there were limits to what could be achieved by natural selection, and that these limits were such as to prohibit natural selection being responsible for the creation of new species.

Section 3.2: What led to natural selection becoming widely accepted?

Various scientific developments eventually allowed the scientific problems facing Darwin’s theory of evolution by natural selection to be dissolved. The most important such developments were in the

study of inheritance and variation, however there were other developments that were also friendly to natural selection.

Advances in understanding of age of earth: I noted earlier that Kelvin's calculations concerning the age of the earth created a problem for evolution by natural selection by suggesting there would not have been enough time for it to take place. The flaw in Kelvin's calculations was eventually discovered by Pierre Curie and Lord Rayleigh in the first decade of the twentieth century (Bowler 1985 p.195). Curie showed that heat was released by the radioactive decay of elements such as radium, and Rayleigh showed that radioactivity inside the earth would offset the cooling described by Kelvin, allowing for the possibility of the earth being much older than Kelvin had estimated. Further, radioactive dating was developed so that by the 1930s the Cambrian period was known to have been more than 500 million years ago. This dissolved the worry that evolution would not have had enough time to occur through natural selection.

Fossil record becoming less unfriendly: Crucial to overturning the perception that the fossil record was unfriendly to natural selection was the work of Simpson (1944). Simpson showed that the available evidence from the fossil record showed that evolution took place "in the irregular and undirected manner predicted by Darwinism [and did not show] a linearity of development more appropriate to Lamarckism or orthogenesis" (Bowler 1985 p.299). He also provided an explanation of the discontinuity of the fossil record that went beyond just the fact that the record was not complete. He did this by explaining why genetic drift would cause substantial evolutionary change "rapidly and in very small populations" and would therefore "be most unlikely to leave fossil evidence" (ibid). This explanation is perfectly compatible with natural selection.

Advances in understanding of inheritance and variation: The most important scientific developments that allowed for acceptance of natural selection as the crucial mechanism of inheritance were in the study of inheritance. Earlier, I pointed to claims that jointly make up the explanation natural selection provides for evolution. Although the truth of each of these claims was not known at the time, each was subsequently vindicated as genetics developed.

First, what led scientists to conclude that variation is heritable? The answer here is that further developments in Mendelian genetics showed that Mendelian genes can give rise to continuous variation. Briggs and Walters explain how Yule was among the first to suggest that "many genes were involved in continuous variation" (Briggs and Walters 1997 p.67). Nilsson-Ehle in 1909 developed this idea by showing that two different genes are involved in determining the colour of wheat chaff (ibid). Larson writes that "Nilsson-Ehle calculated that if ten different genetic factors affect a trait, then sixty thousand variations of it might exist" (Larson 2004 p.171). Demonstrating a

genetic basis for continuous variation showed that such variation could be heritable. Also important in this regard was the work of biometricians such as Pearson and Galton, who showed significant correlations between, for example, the heights of parents and offspring (Briggs and Walters 1997 p.48-59).

Second, what led scientists to conclude that adaptive traits would be able to spread through a population? First, theoretical work was done to better understand the effect of selection on increasing or decreasing gene frequencies. Bowler points to Punnett's (1915) study of mimicry in butterflies "with a table prepared by the mathematician H. T. J. Norton showing how selection would allow a favourable gene to spread within a population" (Bowler 1985 p.293). Fisher (1930) built on this by showing how "If a particular gene conferred an advantage resulting in a faster breeding rate, it was possible to calculate how rapidly its frequency would be increased" (Bowler 1985 p.294). Secondly, studies showed the effects of selection on the distribution of traits happening in practice such as in the case of "industrial melanism in the peppered moth" (ibid). In the course of just over 50 years between 1848 and 1900 the population of this previously light coloured moth became almost entirely dark coloured as a result of the advantage conferred by dark colouring in an environment polluted by soot in the industrial revolution.

Third, what led scientists to believe that these effects could accumulate to a degree where natural selection could explain the arrival of new species? Recall that a crucial piece of evidence leading scientists to think that variation for selection to act on was Johannsen's bean experiments. Over time "further sets of data became available to compare with Johannsen's results" (Briggs and Walters 1997 p.86) and showed that the ineffectiveness of selection in Johannsen's experiments was down to the breeding systems of beans which show habitual self-fertilisation. There have been "many successful selection experiments with outbreeding organisms" (ibid). Developments such as these removed previous reasons to think there would be a limit to what NS could achieve as most people who thought there would be such a limit thought the limit was created by limited variation for selection to act on. The work of Dobzhansky (1937) on isolating mechanisms was also important for deepening the account natural selection could give of speciation (Bowler 1985 p.298).

Section 3.3: What does the case for natural selection tell us about prediction vs. explanation?

How much of what led to the scientific acceptance of natural selection can be described as NPS? Breaking this question down: first, can revised estimates of the age of the earth be seen as NPS for selection theory? Second, can Simpson's work on the fossil record be seen as NPS for selection theory? Third, can advances in genetics be seen as NPS for selection theory?

With regard to the age of the earth, here the answer seems like a clear no. The earth being very old does little to distinguish between different accounts of evolution available at the time. This scientific development was really the removal of a problem rather than a confirmed prediction.

Simpson's arguments about the fossil record also do not seem like NPS for selection theory.

Simpson's arguments were that the fossil record should not count *against* selectionism – that the fossil record is compatible with selection (Bowler 1985). That is not the same as saying that the fossil record confirms predictions made by selectionism and no other theory.

The most important question is whether developments in our understanding of inheritance and variation can be seen as NPS for the theory of natural selection. It might be possible to frame them as such. For example, one might think it a prediction of the theory of natural selection that a continuous range of heritable variation in traits relevant to reproductive success exist in many populations of organisms. The evidence eventually bore out this prediction. That could perhaps be seen as NPS for the theory of natural selection.

There are two things to say in response to this. First, seeing developments in genetics as NPS for selectionism in this way feels like a stretch. Second, as with intermediate fossils as evidence for evolution, it is not clear that the predictive character of this evidence is important (if it exists at all).

The reason why it feels like a stretch to call advances in genetics NPS for selection theory is that the original formulation of the theory of natural selection made no mention of any detailed account of the mechanisms of inheritance. Any prediction the theory makes about the mechanisms of inheritance has to be understood at quite a high level of abstraction and it seems possible in theory that the needs of natural selection might also have been compatible with other ways that our understanding of the mechanisms of inheritance might have developed. The crucial thing about, for example, the realisation that multiple genes for the same trait can give rise to continuous variation, was that it removed difficulties that had been suggested for natural selection by earlier ideas of inheritance and variation.

The reason why any predictive character that might be read into this evidence feels of limited importance is that there does not seem to be anything plausible to say about why the evidential situation would be greatly different had a full account of the mechanisms of inheritance and variation happened to precede the suggestion that natural selection drives evolutionary change. The degree to which the understanding of genetics that prevailed in the early 20th century is amenable to synthesis with natural selection would be the same in that case as it was in the real history of science, where the order was reversed. Even if we can convince ourselves that genetic

developments were NPS for natural selection, the status of these developments as NPS does not seem to change their epistemic importance in any significant way.

Conclusions

There are two conclusions to these case studies I want to push. First, whether the motivations for a scientific theory can be massaged as NPS or not is of limited relevance for the plausibility of realism about that theory. I take it that realism about the theory of evolution is as plausible as realism about any scientific theory. However, the confirmation of predictions made by this theory seems to have little role to play in explaining why our grounds for the theory were so powerful in the late nineteenth century. Similar points also apply to the theory of natural selection. This undermines attempts to phrase an STI that is supposed to apply throughout science only in terms of NPS.

Second, the actual reasons for adopting a theory in science are often diverse enough that seeing them all as instances of one type of evidence is difficult. In the cases above, we have seen that trying to apply the concept of NPS to much of the evidence for evolution and natural selection seems to misrepresent the reasons why that evidence is powerful. This gives some early indication of some points I will be developing later in this thesis. These are that, first that context is important in determining the strength of different kinds of evidence. For example, the fact that evolution is one of only two possible accounts of the origin of species changes that force of evidence for the theory considerably. Second a range of kinds of evidence are often in play in any scientific situation. These different kinds of evidence often cannot easily be forced into a single category without that category losing any important role in explaining the force of those different kinds of evidence. In the remaining chapters of this thesis I will start to work towards an alternative framing of realism debate that better accounts for this diversity in science.

Chapter 4: A Different Approach to the Realism Debate

Introduction

In the previous chapters I outlined how the narrow understanding of success that realists have been working with in the contemporary debate over the success to truth inference (STI) is unable to properly account for the rationality of some very good science, and is poorly motivated anyway. This leaves the realist in a difficult situation, as shifting from a general notion of success to the notion of novel predictive success has generally been viewed as one of the key moves in the realist's defence against the pessimistic induction. If this move fails, what is the realist left with?

My answer is not that the realist is left in a hopeless position and the anti-realist in a favourable position, but that there is something wrong with the terms of the whole debate. The problem is that the STI is presented as the single criterion for deciding whether *any* scientific theory whatsoever is worthy of realist commitment. However, no single inferential rule is capable of playing that role. The methods of science involve many different inferences, values, rules, and principles, each of which needs to be interpreted in the light of background beliefs about the world that can vary, and each of which draw at least some of their force from the context in which they are used. This diversity in the methods and topics of science means that any attempt to capture the grounds for realist commitment to any theory in science using a single criterion will not succeed. The answer is not to abandon the possibility of realism, but to recognise that the grounds for realism vary. Rather than doing away with the STI or PI, we should have many of each instead of one of each.

I will begin my case for this position by arguing that the realism debate should be interpreted as a methodological debate. The STI should be seen as a methodological principle intended to guide theory choice. Historical arguments such as the PI should be seen as part of the project of methodological naturalism – i.e., the idea that methods can be tested empirically by examination of their track record. This kind of position is frequently found in discussions of meta-methodology in the scientific method literature (Sankey and Nola 2007). Notable advocates include Quine (1992), Laudan (e.g. 1996) and Rescher (1977).

Once we situate the scientific realism debate within the literature on scientific method, a number of new questions present themselves. If the STI is a methodological principle, does it satisfy the general requirements philosophers have suggested for methodological principles in science? In what ways are methodological principles amenable to empirical testing, and what problems stand in the way of this means of justifying methodological principles? Is 'truth' a transcendent property of theories, and if so, is it thereby unsuitable for use in methodological principles? How well does the STI represent

the methods actually used by scientists? If the answer is that it represents them poorly, what are the implications of that? What implications does the widespread acceptance of methodological pluralism have for the scientific realism debate? In what ways can historical assessments of methods go wrong?

My most important aim in this chapter is to show that the global phrasings of both the STI and the PI are committed to a view of scientific method that is no longer at all popular, and rightly so. This is the view that there is a universal scientific method that applies throughout science. Once I have established this, I will go on to present the arguments in favour of methodological pluralism, and therefore against the traditional formulations of the STI and PI, in chapter 5.

Section 1: The Realism Debate as a Methodological Debate

The STI sets out a method of forming beliefs, namely ‘inferring (probable approximate) truth (of a theory’s working posits) from (the right kind of) success’, and claims that this method is reliable for the topic of ‘science’. This seems directly analogous to any other context where we claim that a method of forming beliefs is reliable. For example, we might make the same sorts of claims when evaluating testimony: ‘I am justified in believing that it will rain tomorrow because the Met office say so’. This claim implies that the Met office are a reliable source of information about tomorrow’s weather, and that forming my beliefs according to what they say represents a reliable method of forming beliefs.

Similarly, the PI argues that a particular method of forming beliefs in science, namely inferring truth from success, is not a reliable method for the topic of science by attempting to show that following this method has not always led to true theories.

Once we conceive of the STI and the PI as arguments about the reliability of methods, we have a new way to look at the realism debate. By making it explicitly a debate about the reliability or otherwise of a method of forming beliefs we are able to take insights from the philosophical literature on scientific method. This turns out to have important consequences for how we engage in the realism debate.

Sankey and Nola (2007) label the task of justifying a scientific methodology as ‘meta-methodology’. They present a three tiered analysis of the topic of scientific method. At level one are particular scientific judgments, for example, the choices of theories that scientists have settled on in different areas. Theory choices at level one are guided by methodological principles at level 2. These principles include “deductive and non-deductive principles of inference; values by which theory choices can be made; and methodological M-principles, which govern theory choice and can include

values.” (ibid p.80.) M-principles is short for methodological principles, and can take the form of categorical imperatives (i.e., ‘you should do x’) or hypothetical imperatives (i.e., ‘you should do x if you want to achieve aim/ value y’). Sankey and Nola give as examples of M-principles that have been endorsed by theorists of scientific method, “One ought to avoid ad hoc theories” and “given some rival hypotheses and a large set of facts they all explain, then one ought to pick that hypothesis that best explains all the facts” (ibid p.59). The STI can be expressed as an M-principle as follows: “if a theory shows the right kind and degree of success, infer that it is probably approximately true in its working posits”.

Level 3 of Sankey and Nola’s scheme is meta-methodology, in which the principles that feature at level 2 are justified. Here we explain why it is right to follow one possible methodology rather than another. Sankey and Nola outline three ways philosophers can or have attempted to justify methodological principles. First, we can offer *a priori* arguments in favour of a methodological principle. Second, we can attempt to offer an empirical justification. Third, we can attempt to justify our methodological principle using the method of reflective equilibrium, in which justification is obtained for general principles by showing that they are in alignment with our particular scientific judgements, and that general principles and particular judgments are mutually supporting.

We need only focus on the second method mentioned by Sankey and Nola, i.e., using empirical arguments in favour of a methodology. To illustrate what Sankey and Nola have in mind here remember that M-principles can be expressed as hypothetical imperatives – ‘if you want to achieve aim/ value y, do x’. A hypothetical imperative of this kind is grounded by the causal claim, ‘doing x leads to y’. This claim is amenable to empirical test. We can look at all the cases where x has been done, and see how often it led to y. If x usually did lead to y, this provides justification for the M-principle. If it did not lead to y as often as we would expect, then this undermines the M-principle.

Applying this to the realism debate, we can phrase the STI as a hypothetical imperative as follows: ‘if you want theories that are true at least in their working posits, accept only those theories that show an appropriate kind and degree of success’. This is grounded by the claim ‘theories that show an appropriate kind and degree of success are probably approximately true in their working posits’. This is then amenable to the kind of historical tests that advocates of the pessimistic induction have attempted to subject it to. Have most appropriately successful theories turned out to be approximately true in their working posits (at least by our current lights), or not? Anti-realists have been concerned with showing the answer is negative, and realists with resisting that conclusion.

Translating the realism debate into a meta-methodological question is barely even a translation. What is going on in the realism debate is fairly clearly an empirical test of a methodological principle.

However, the methodological nature of the realism debate is not often explicitly acknowledged. This means that the insights of the scientific methods debate are sometimes overlooked, even though they are highly relevant.

At this point someone might object that the STI is not supposed to be a methodological principle adopted by *scientists*, but a tool used by philosophers considering science from an outsiders perspective. There are two problems with this idea. First, the STI is supposed to be a way of deciding which theories we are justified in believing are true. That is very much a concern of scientists as well as philosophers, so if it is a good way of deciding which theories are true, then why would scientists not want to adopt it into their methodology? Second, evaluating the degree of success shown by a theory is not something that can be done from 'outside' science anyway, since it requires a detailed engagement with the first-order evidence for that theory. Once they got themselves into a position to legitimately judge the degree of success shown by a theory, the philosopher would no longer be outside science but would be acting as a scientist. I therefore see no reason why the STI should not be evaluated as an account of scientific method.

However, perhaps the objection expressed above is best put another way. The STI is a rule about when we should regard scientific theories as literally true, whereas the rules of scientific method are rules about when we should adopt theories. 'Adopt' does not necessarily mean the same thing in science as 'believe to be true'. There are other senses in which a scientist might adopt a theory. For example, a theory could be accepted as a working hypothesis, or as empirically adequate (see van Fraassen 1980), or as a tentative finding, or as an established fact. The last of these matches up with the philosophical notion of realist commitment, but the others do not.

This objection perhaps relies on the idea that viewing the STI as a methodological rule means that we have to see truth as the overarching goal of science, and that this is not legitimate. The idea that truth is such a goal has been disputed by philosophers such as Laudan (1984) and Van Fraassen (1980), who argue instead that science is better conceived as having some other aim, such as empirical adequacy in the case of Van Fraassen. If viewing the STI as a methodological rule *does* rely on conceiving of the aim of science in this way, then that may create a way for this way of seeing the STI to be contentious.

However, we do not necessarily need to see truth as the overriding goal of science in order to think of the STI as a methodological rule. M-principles can be expressed as hypothetical imperatives, 'if you want to achieve value y, do x'. Even if 'truth' is just one among a variety of values in science, then the STI can still be a methodological rule, albeit a conditional one that only applies when we are interested in our theories being true.

In any event, that science aims at truth is itself a key commitment of scientific realism so if we accept the STI in any capacity we should be comfortable with accepting it as a methodological rule of science.

The point about there being different ways in which a scientist might adopt a theory is correct, but of limited importance. We can just think of the STI as a methodological rule that applies specifically and only in scientific contexts where we are deciding whether a theory or claim should be adopted in the sense in which adoption implies believing to be literally true.

Section 2: Problems for Testing Methodological Principles Empirically

The use of the historical track record of a methodological principle as a means of testing that principle empirically is an idea that has come up before in scientific method literature. Quine (1992), Rescher (1977), and Laudan (1996) all recommend this approach, which is known as normative naturalism or methodological naturalism (Sankey and Nola 2007).

Two main problems face empirical meta-methodological techniques. First, a problem is posed by the question of what methods we are to use in our historical evaluation of a method. If the same methods we are meant to be evaluating are used in the evaluation, we run into a problem of circularity. If the methods we are meant to be evaluating are not used in the evaluation, then we have to ask what justifies *those* methods. This leads into a regress problem. Although an important point, this is a general problem for justification, not just for the study of scientific methods. I will not explore this issue in any great detail here.

A second problem facing empirical meta-methodology is raised by confirmatory holism. Confirmatory holism is the thesis that our beliefs face the test of experience 'as a whole'. Say we want to test a theory. That theory makes certain predictions which we set about trying to verify or falsify. Say that one or more of these predictions turns out to be false. This situation calls for some revision in our beliefs. However, we have a choice about where to make this revision. No theory makes predictions on its own. It requires various supplementary hypotheses such as 'the scientific equipment is working properly', or, 'the sample chosen is representative of the population' and so on. Often it is these supplementary hypotheses we will choose to revise rather than the theory itself.

Applying this to empirical meta-methodology, if a rule seems to be refuted by the history of science this does not necessarily mean we need to immediately reject the rule. If a rule takes the form 'do x if you want y' we can try to evaluate it by looking at the track record of x being done to see whether it achieved y. Say that we find cases where doing x did not achieve y. What should we conclude? It might be that the rule we are testing is no good, or it could be one of a number of other things.

Perhaps the rule was not properly applied in some or all of those cases. Or perhaps in some cases another methodological rule was in conflict with the one we are trying to test, and priority was given to the wrong rule.

Methodological principles do not function in isolation, and do not by themselves lead to any outcome. Just knowing that you are supposed to prefer simpler theories for example does not get you very far. Methodological principles function within an overall methodology that includes many other principles, inferential rules, values and so on. Methods are used in a context featuring background beliefs about the work, particular research puzzles we are trying to solve, and so on. If we find a case where some scientist applied a specific methodological rule but this did not lead to the outcome we would have expected, it may be because they did not properly appreciate how to balance this rule against others. An understanding of how such balancing *should* be done can legitimately be seen as part of a methodology. If we see a method in this sense as the thing we are evaluating by comparison with the historical record, then the problem of holism starts to become clear. When we find some case where the method was applied and the outcome was not as expected, it may be any of the different principles that make up the overall method, or the balancing rules of those principles, that is to blame.

In the context of the scientific realism debate, these points may seem irrelevant, since the method under discussion has only one methodological principle, i.e., the STI. However, once we see the STI as an account of scientific method, it becomes obvious that some level of further detail is likely to be needed to make it workable. Some account will be needed of how all the various kinds of scientific evidence contribute to success. Some account will be needed of what different kinds of success there are, and what relative effect each has on overall success, and so on. As a result, confirming or disconfirming the STI inevitably becomes more complicated. In fact, at this point it begins to look as if the STI can only be at best a gesture in the direction of an account of scientific method, rather than an account in itself.

However, if that is the case then it points to a general problem with how the realism debate has been conducted. On the basis of the STI and the PI philosophers have sought to draw conclusions about the epistemic status of science in general. The STI boils the case in favour of scientific theories down to the 'success' they show, and in doing so is doing one of two things. If 'success' is a given a single, reasonably definite meaning such as NPS then it is discussing a method of forming beliefs that is of relatively low importance in many areas of science, and therefore not a good way of deciding the viability of realist commitment in those areas (see chapter 3). If 'success' is allowed to stand as a more general placeholder for various kinds of scientific evidence then the STI can only stand as a

placeholder for a much more detailed theory of scientific method that somehow ties together each of these different kinds of evidence, with 'success' somehow providing a unifying underlying logic to each of them.

Returning to the theme of confirmatory holism as a problem for testing methodologies empirically, if we see the thing to be tested as a whole methodology rather than individual methodological principles, and a methodology comprises many principles together with balancing principles, inferential rules and so on, then we are trying to choose between an almost infinite set of possible methodologies. This presents a problem because at least some degree of methodological change is likely to have happened regularly over the history of science and in that case there may not *be* a significant track record for any particular methodology, if we conceive a methodology as an exact set of principles and balancing rules etc. This general theme will recur over the next few chapters – individuating methods in science to find the proper unit for empirical testing against the history of science is not easy. This will present a major problem for finding the right level for localising the STI and the PI in chapters 5 and 6, where I will discuss the issue in detail.

Section 3: Can 'Truth' Feature in a Testable Methodological Principle

Another potential problem with seeing the realism debate as a methodological debate is discussed by Sankey (2000). This problem is that, according to Laudan (1996), 'truth' is a transcendent property and therefore unsuited to acting as a goal or value for science. By 'transcendent', Laudan means that we do not have access to whether a claim is true or false – there is no way of proving empirically beyond all doubt that a claim or theory is true. At best we can prove a claim or theory false. As such, we cannot know when truth has been achieved, making it unsuitable to be a goal of science. However, if it is unsuitable to be a goal then it is also unsuitable to feature in a methodological criterion such as the STI because any claim of the form 'doing such and such leads to true theories' is impossible to confirm empirically. We do not have access to whether any theory *is* true – at best we can just know that it hasn't yet been shown to be false. This means we have no way of evaluating the success of any method meant to achieve truth as a goal, and so there is no way for us to evaluate the STI as a methodological principle by its track record.

Sankey objects to Laudan's claims on the basis that our inability to know infallibly whether our theories are true is not the same as us having no access to truth whatsoever: "It may readily be conceded that there are no infallible criteria of truth. But it by no means follows that there are no fallible criteria for the recognition of truth." (Sankey 2000 p.223.) With regards to whether a method is conducive to truth Sankey says, "It is perhaps true that there may be no evidence that a method leads to theoretical truth. But there may surely be indirect evidence that a method conduces to such

truth" (ibid). The sort of indirect evidence that Sankey has in mind is based on abductive reasoning. He argues that the best explanation of the success of science is that the methods of science are truth conducive.

Applying this to the current project of seeing the STI as a scientific method, this argument of Sankey's has a whiff of circularity about it. However, the basic point that our lack of infallible access to truth need not make it unsuitable as long as we have access of some kind seems sound. The weight of argument behind many scientific theories strikes those unfamiliar with the PI as good, albeit perhaps not fool-proof, reason to think those theories are true. This reasoning may or may not be shown to be incorrect by historical arguments like the PI, but so long as PI type arguments can be resisted, it seems reasonable for the idea that success is a guide to truth to be a default. If that is the case, then if we look into the history of science and find that theories showing the appropriate sort of success all turned out to be approximately true in their working posits there seems no reason not to take that as strong but not infallible evidence for the reliability of success as a guide to truth. Methodological principles featuring truth as an aim should not be regarded as inherently incapable of empirical testing.

Section 4: Implications of Seeing the Realism Debate as a Methodological Debate

What are the implications of treating the realism debate as a methodological dispute? First, it prompts the question of how historical assessments of methods can go wrong, and whether the STI or the PI go wrong in those ways. Second, it raises the question of how well the STI reflects actual scientific methods. Third, it raises the question of how scientific realism intersects with methodological pluralism. The STI is supposed to tell us the method for deciding whether to believe any scientific theory is true. However if methodological pluralism is correct then there is no one method that is used across science. This suggests that methodological pluralism is incompatible with a global approach to the STI or the PI.

I detail each of these implications below.

Section 4.1: Two Possible Errors for Historical Assessments of Methods

The pessimistic induction is an instance of a type of argument that could be used against targets other than science in epistemology. The abstract argument type that the pessimistic induction is an instance of is as follows:

1. In the past method A has led us into error
2. There are no relevant differences between the situations where A led us into error and our current situation

3. Therefore, we do not know that A is not leading us into error now

The argument begins by noting a method's history of being unreliable, goes on to claim that we cannot distinguish past situations where we were in error from our current situation, and concludes that we cannot take the method we are considering to be a reliable route to knowledge in the situation we are in. Let us call this general argument type the inference from past unreliability.

In the pessimistic induction, method A is filled in with the realist's latest formulation of the general 'inferring truth from success' motif. However, there are other instances of this argument where A is filled in some other way.

The most famous example of this type of argument is Descartes' dream argument (Descartes 1998 [1641] First Meditation). There Descartes notes that in the past what he took for perceptual experiences were actually part of a dream. When he was dreaming, his experiences seemed to him to be genuine sensory perceptions. He concluded that he did not know that he was not dreaming at the time he was writing, and therefore did not know the things his senses seemed to be telling him. This argument can be recreated by filling in method A with 'what seems like sense experience' in the argument structure above.

A more everyday example of this kind of argument is in our evaluation of people's testimony. When we are deciding whether or not to trust what someone tells us, we often try to apply the argument above and see if it sticks. We assess the track record of the individual to see whether they have generally been reliable in the past. If they have not been, and we have no way to distinguish the situation where they were unreliable in the past from the situation now, we may take them to also be unreliable now.

Considering the argument from past unreliability as it applies to people's testimonies can help to illustrate where the argument can go wrong. Two major ways the argument can go wrong are:

1. It can be applied in a way that fails to recognise differences in the reliability of the same methods across different contexts.
2. It can be applied to multiple methods masquerading as a single method

I illustrate each of these possible errors further below.

4.1.1: Conflating Topics

To illustrate the first way the argument from past unreliability can go wrong, consider a person, Megan, who is a grandmaster at chess but has a terrible sense of direction. In the past Megan has

tried to give me directions and got me horribly lost. This has happened on multiple occasions. Now Megan is trying to explain what went wrong in a chess game I lost. Consider this argument:

1. In the past Megan has often given me wrong information (*this is true because she has got me lost so many times*)
2. Therefore I should not trust what Megan is telling me about my chess game

This argument is silly because the same method of forming beliefs (in this case Megan's testimony) can be reliable on some issues and unreliable on others. This instance of the argument from past error to general unreliability fails because it tries to move from errors on one topic to unreliability on another without any good reason to think that the reliability of a method of knowledge acquisition will be similar for both topics.

Whether an argument goes wrong in this way is not so much a matter of counting the topics at play in an argument, as deciding whether the topics at play are appropriately related. Take the following example. Tom is someone who got a poor grade on his GCSE maths exam and has not studied maths any further since. However, he is also someone who likes to appear knowledgeable on everything and will sometimes feign knowledge he doesn't have. In the past I have asked his help with algebra questions I couldn't answer and he has given me what I later found was wrong information. Now he is offering help with a geometry question I am struggling with. Consider the following argument:

1. In the past Tom has given me wrong information about algebra
2. Therefore I should not trust what Tom is telling me about my geometry problem.

This example of the argument from past unreliability seems okay. Even though it does move between topics, the topics it moves between are related in such a way as to let us think we can move from premises about one to conclusions about the other in this instance. Both fall under the general heading of mathematics and we have good reason to think that Tom is likely to be equally unreliable on all mathematical questions since we know he has not studied it to any high level.

Another thing to consider when we are wondering whether an argument from past unreliability fails because it conflates topics is that sometimes we might think that a candidate route to knowledge is so suspect that we can reasonably think it unreliable on *all* topics. If we think someone is a pathological liar then we can reasonably distrust what they say about most things, without necessarily needing to consider the specifics of things like their expertise on the subject they are talking about.

So noting that an argument from past unreliability conflates topics does not automatically show that argument to be problematic. It does however raise the question of why we should think the method

of belief formation we are considering is likely to have the same degree of reliability across the two topics. Equally, one might ask why it should have a different degree of reliability across the two contexts. Where the burden of proof falls here is likely to be complex and context dependent – it will depend on what our background beliefs are about both the method we are considering and the topics we are considering. These are all points to bear in mind as we discuss whether the pessimistic induction falls to the same kinds of problems the argument about Megan above displays.

4.1.2: Conflating Methods

To illustrate the second possible pitfall in historical assessments of methods of forming belief, imagine we have two people, John and Mary. Mary is a paragon of epistemic virtue and honesty, while John spends no effort finding out what the truth really is about anything, and even when he knows the truth he usually lies about it. Imagine we now treat the testimony of John and Mary as a single candidate-route-to-knowledge in the following argument:

1. (John and Mary) has told me that the shop is closed. *(This is true because Mary has told me the shop is closed.)*
2. In the past (John and Mary)'s testimony has often been false. *(This is true because John has often said false things and John is part of (John and Mary).)*
3. Therefore, I do not now know on the basis of (John and Mary)'s testimony that the shop is closed.

Although the argument above may have the surface form of a reasonable inference, there is clearly something suspicious about it. There is no reason to treat (John and Mary) as the target of this argument because John and Mary differ widely in their reliability. Even if they didn't they would still represent different possible paths to knowledge. There is no general reason to think that what is true of John's reliability will also be true of Mary's reliability when it comes to whether the shop is open. Nothing is gained by talking about John at all in the argument above.

Of course, no one would be tempted to make the exact mistake seen above. However, more realistic situations can be devised in which such an error would be easier to make. We often take different people or groups to be associated in a way that makes us inclined to treat them the same way epistemically, even though we ought not to. For example, in the UK a few years ago some politicians and media commentators expressed a disinclination to trust 'experts'. Part of the motivation for this may have been a perception that what 'experts' had said in the past had often turned out to be wrong. However, the reliability of different people who could all be labelled 'experts' is largely

unrelated. So in this instance the argument from past unreliability would not have been well applied because it would have been similar to the John and Mary situation above.

Whether or not an instance of the argument from past unreliability can be convicted of this kind of mistake really depends on whether we have reason to think that the sources of information we are discussing are related in a way that makes them likely to be equally reliable or unreliable with regard to the topic under consideration. In some cases, we could be justified in treating multiple people's testimonies as collectively making up *one* candidate route to knowledge. For example, say we hear what one government minister has to say about a subject, then the next day we hear another member of the same government talk about the same subject. Do we have one candidate route to knowledge here or two? Plausibly the answer is one, since it is quite a safe assumption that the members of the government have discussed the issue together and decided on the 'party line' to be presented in media interviews by all members of the government. Little independent corroboration is provided for one minister's statement by another's.

If we were to discover that an argument we endorsed appeared to treat multiple methods of forming beliefs as if they were one there would be two choices. First would be to argue that the different methods we had been treating as one are sufficiently closely related to make it justifiable to act as if they *were* one for the purposes of the argument we are making. The second option would be to make our argument more specific by rephrasing it only in terms of one of the candidate routes to knowledge we had previously been conflating – this is generally likely to make the argument more solid but also narrower in scope.

Section 4.1.3: What these types of errors tell us about the STI and the PI

Applying the observations above to the PI and the STI, if either of these arguments commit either of these errors then the other does as well, because both take the same method and the same topic as their target, viz. 'inferring truth (approximately in working posits) from (right sort) of success', and 'science'.

This tells us that if 'inferring truth (approximately in working posits) from (right sort of) success' turns out not to be one method, but rather to be multiple methods in disguise then the STI and PI are guilty of conflating methods. If 'science' turns out not to be one topic but a collection of not relevantly related topics, then both the PI and the STI are guilty of conflating topics. In either of these situations it will be necessary to narrow the scope of the PI and the STI to address only one method and one topic, unless there is a good argument in favour of lumping together separate methods or topics.

In subsequent chapters I will be arguing that traditional arguments in the realism debate do commit both of these errors, and that we should therefore replace the versions of these arguments that are supposed to apply to all of science with multiple localised arguments meant to apply to only one method and one topic. Identifying how to individuate methods and topics will be the major challenge facing this project.

Section 4.2: How Well does the STI Reflect Actual Scientific Methods?

Thinking of the STI as a methodological principle also invites the question of how well the STI represents the methods actually used by scientists. When scientists are evaluating the case for a theory do they, or should they, consider the effect of any given argument or piece of evidence on the overall amount of success that theory shows. The answer to this is likely to depend to at least some degree on what concept of 'success' we are working with. If it is novel predictive success, then I have already argued at length in chapter 2 that the answer is often 'no'. But perhaps that just shows that some other concept of success is needed in the realism debate. Perhaps we just need to re-broaden the concept of success to include explanatory success for instance.

The problem with moving to a broader concept of 'success' is that either it remains too specific to take in much of science, as we saw with NPS based concepts of success, or it becomes rather vague. Explanations of what explanatory success *is* usually tie together a number of different features that good explanations might have in different circumstances, such as breadth, simplicity, and so on. The value of each one of these values might itself be the subject of a methodological principle, but it remains to explain what is added by the aggregative concept of 'explanatory success'. In order for 'explanatory success' to provide any meaningful criterion of realist commitment we would need to have a way to judge explanatory success in a way that took account of the balancing act often needed between different desiderata for explanations. Under which circumstances should we choose the simpler explanation, and under which circumstances should we choose the broader explanation? This is complicated even further by the fact that there might be different answers to these balancing questions, and by the fact that the answer might be sensitive to our particular context. This raises the threat that explanatory success might mean one thing at some times and another at others.

The idea that what 'success' means is not always the same is extremely problematic for the traditional STI and PI however. If 'success' does not always mean the same thing then there is a serious risk that both the STI and PI are making the error of conflating methods, since both present some version of the basic 'infer truth from success' idea as *one* method of belief formation to be evaluated. If 'success' does not always mean the same thing though, there is a separate method for

each of the different meanings of 'success'. If that is the case, then the onus is on the advocate of the STI or the PI to explain what unifies the different meanings of 'success' in a way that makes them amenable to collective evaluation.

This sort of point doesn't just apply to explanatory success, but to any kind of success concept. Generality is bought at the cost of vagueness that it is hard to suspect disguises a lack of uniformity in meaning across science. What we need is for our success concept to in some way tie different sorts of scientific evidence into one overarching notion of success that creates a sufficiently strong epistemic link between each of these sorts of evidence to make evaluating them all collectively as instances of 'success' make sense. What we are looking for ideally is a uniformity of method across the whole of science in terms of the evaluative standards used by scientists. If there is one general method of science which sets out how to evaluate any kind of scientific evidence against agreed standards, then these evaluative standards can give a content to the notion of 'success' that allows 'infer truth from success' to be evaluated as one method and which also allows the STI to reasonably be considered representative of the whole of science. The kind of view of scientific method that this would rely on is known as universalism, since it holds that there is a universal scientific method. However, this idea has come under sustained attack over the last fifty or so years from methodological pluralists, as I explain further in the next section.

Section 4.3: Methodological Universalism, Methodological Pluralism, and their Implications for the Realism Debate

Universalism about scientific method is the idea that there is a single universal scientific method that is used throughout science. The use of this method is thought of as constitutive of an activity counting as science. Methodological pluralism is the view opposed to this general way of understanding scientific method. Methodological pluralism is summarised by Sankey (2000) by the following five points:

- "Multiple rules: scientists utilise a variety of methodological rules in the evaluation of theories and in rational choice between alternative theories.
- Methodological variation: the methodological rules utilised by scientists undergo change and revision in the advance of science
- Conflict of rules: there may be conflict between different methodological rules in application to particular theories.
- Defeasibility: the methodological rules, taken individually rather than as a whole, are defeasible.

- Non-algorithmic rationality: rational choice between theories is not governed by an algorithmic decision procedure which selects a unique theory from among a pool of competing theories.” (Sankey 2000 p.213.)

These five claims need not always go together, however Sankey regards them, correctly in my view, as capturing the most central points of pluralism as a general approach to methodology. As Sankey phrases them, the first two of these claims appear descriptive, but the approach of scientists they describe is also generally given normative support by pluralists. That is to say, pluralists hold that there is something either inevitable or desirable or both about the existence of multiple methodological rules in science, and the variability of methods across fields of inquiry and moments in science’s development.

I will leave the arguments for and against universalism and pluralism about method to the next chapter. In the remainder of this chapter I want to finish arguing for the conditional claim that global formulations of the STI and PI can be defended *only if* universalism about method can also be defended.

The STI and the PI are usually used to license conclusions that are supposed to apply to science in general. The intended conclusion of the STI is that (the working posits of) any scientific theory whatsoever is a good candidate for realist commitment if it has shown appropriate success. The intended conclusion of the PI is that success is not a reliable guide to truth in any area of science. If we interpret these as methodological claims then they rely on the assumption that there is one universal scientific method which applies throughout all of science. This is because the STI suggests itself as a methodological rule that applies throughout all of science, and the PI attacks it on those terms.

However, if methodological pluralism is true then science is made up of multiple distinct methodologies that do not reduce to any one method, and are not meaningfully united by any one method. If that is the case, then the STI can only apply to some of science at best, and even if the PI undermines the STI successfully it does not thereby license anti-realism about scientific theories in general.

Chapter 5: Methodological Pluralism and Individuating Methods for Historical Assessment

Introduction

In the previous chapter I argued that the traditional framing of the scientific realism debate can be justified only if universalism about scientific method can be supported. In this chapter, I will explain why this state of affairs is unfortunate for the traditional framing of the scientific realism debate. Very few people would still regard universalism about scientific method as a tenable position. The extent of the consensus on methodological pluralism is such that scientific method is regarded as 'yesteryear's debate' according to Sankey and Nola (2000), since the lack of any overarching scientific method has led many to conclude that little of much generality and interest can be said about scientific method.

Throughout this chapter a repeating theme will be that when talking about scientific methods, there is always a pressure to frame methods in a narrower way. The more generally applicable a scientific method is supposed to be, the 'thinner' it has to be in order to accommodate the diversity of specific methods found in science. The most extreme example of a generally applicable scientific method is obviously one that is supposed to be universal to science. The main reason why universal accounts of scientific method seem unworkable is that even the most abstract and basic points about science appear not to apply to *all* of science. Furthermore, even if we ignore such exceptions, accounts of scientific method vague enough to be so universally applicable are too vague to be put to much use.

My hope however is that a version of the historical scientific realism debate will still be salvageable. Once we admit that there is no unified method in science that may or may not lead reliably to true theories, the natural step is to try to empirically test each of the more specific methods of science against the history of science to see if they reliably lead to true theories or not.

As we will see however, identifying and individuating methods in a way that allows for this is not easy. There is always pressure to move to a narrower and narrower conception of methods, which in the limit ends with the idea that methods are so highly specific that each theory is formed and justified by a different method. This idea is incompatible with the possibility of arguing successfully that any abstractly considered methods are reliable or unreliable.

After arguing against universalism about scientific method at the beginning of this chapter, I will go on to outline two possible motivations for pluralism about method: first the existence of multiple systems of methodological rules; second the contextuality of methodological rules. As examples of the first sort of pluralism I will consider the work of Kuhn and Hacking. As examples of the second

sort I will look at Feyerabend and Norton. I will argue that both sorts of pluralism about method are hard to reconcile with even a localised version of the historically focused scientific realism debate. However, multiple systems of rules type pluralism does perhaps allow some limited scope for historical inductions to establish the reliability of a method.

Section 1: Universalism and its Problems

Providing a general theory of scientific method was a widespread pre-occupation of philosophers of science in the first half of the twentieth century.

There were a number of reasons for this interest. Philosophers were concerned to give an account of scientific rationality. That is, they wanted to explain what determines theory choice and why this demonstrates the superior rationality and reliability of science as compared to non-science. The point about the differences between science and non-science was also very important – the question of how to distinguish the two is known as the ‘demarcation problem’ and was a central issue for Popper (1959, 1963).

Giving an answer to what makes ‘science’ rational, or what makes science different to non-science requires assuming that science is a reasonably unitary thing. To distinguish sharply between science and non-science we need there to be a characteristic method of science. Then, if an activity follows that method it is scientific and if it does not then it is not.

This already gives a hint into the sorts of arguments that led to universalism about scientific method ceasing to be popular. As a general position, it is very exposed to any argument that can show an example of respectable science failing to follow the method prescribed by a universal theory of method. Since a universal theory is supposed to be universal, any such cases pose a serious threat.

Around the middle of the twentieth century a movement took hold in philosophy of science known as the ‘historical turn’ (Sankey and Nola 2007) in which philosophers emphasised the need for theories of scientific method to be informed by the actual history of science. Members of this movement presented many arguments in which historical scientists appeared to use methods that do not conform well to the doctrines of the theorists of method who came before them. Many of these historicist writers presented these cases as arguments against the possibility of a universal method of science, since the methods found in the history of science were too diverse to make their being captured by a single methodology plausible. I will outline the arguments of some members of this historicist turn later in the chapter.

An illuminating passage of *The Disorder of Things* by John Dupre (1993), illustrates how the diversity of science makes it challenging to say anything about scientific method that is applicable throughout

science. Dupre presents three claims that he suggests might be taken as providing a means of differentiating science and non-science. First, science features an “essential connection with dependence on empirical evidence” (1993 p.239). Second, “we expect science to be cooperative and cumulative” (ibid). Third, “we might expect a genuine science to include, or at least seek, laws or principles of considerable generality” (ibid).

Dupre argues that it is implausible that any collection of such principles be “jointly sufficient or individually necessary for a human practice to count as a science” (ibid). He does this by arguing that for each of the features of science he mentions, there are cases of science not showing those features. These aspects of science seem so fundamental to it that their failure to be universal appears to seriously undermine the viability of any attempt to set out the methods or characteristics that are fundamental to science as an activity.

Against the link between science and empirical evidence Dupre says that the link between theory and empirical evidence “becomes much murkier with those areas of science that attempt to provide theoretical models of intelligible simplicity for admittedly enormously complex processes. I have in mind, for example, population genetics, mathematical ecology, or most of economics.” (1993 p.239.) Because they are about such complex systems, it is very hard for their theories to be refuted or confirmed by empirical evidence. By contrast, Dupre notes that “history, aerofoil design, and violinmaking are highly empirical activities” (ibid) despite not being generally regarded as sciences.

Dupre goes on to make similar points about the other two features of science he mentions. He claims that some of science is not co-operative, and cites Mayr’s introduction to *Birds of the Southwest Pacific* in which Mayr notes that his is the only work on the subject in English. Dupre also argues that some of science is not cumulative. He argues that “more-theoretical debates... seem to circle around inconclusively in the manner often attributed to traditional philosophical problems” (1993 p.240).

Finally “the role and status of [broad theoretical principles] will vary greatly from one area of scientific inquiry to the next” (p.241). Examples of areas where they are of little importance include human psychology and sociology for Dupre.

Dupre’s arguments illustrate well the sort of challenge facing any attempt to give a single universal account of the methods of science or generally how science works. Science is very variable. If we go looking, it is likely to be easy to find areas of science that look like outliers in terms of their methods, goals or other characteristics. Most relevantly for us, the particular methods used in different areas of science have a surface appearance at least of great variability. In terms of the equipment

scientists use, the practices they engage in, the methods of observation they use, and more, diversity is what is most striking.

It is very challenging to see what *could* feature in a universal theory of scientific method. Some of the aspects of science discussed by Dupre seem both extremely vague and general, and also central to our understanding of how science works. If even these sorts of principles are not universally applicable in science then it is hard to see what might be. We will see further support later for this line of reasoning in the work of Kuhn and Feyerabend. Kuhn argued that changes in fundamental theory cannot be made on entirely logical or empirical grounds, and Feyerabend argued that even such principles as 'do not adopt *ad hoc* theories' cannot be universally applied. When even the most general and basic attempts to formulate principles of universal scientific method turn out to be faced with such serious exceptions, the prospects for any kind of universalism about scientific method begin to feel bleak.

As the discussion of universalism above indicates, the case against universalism often arises from the details of pluralists' historical analyses. In the next section I will outline some of the accounts of scientific method put forward by some influential pluralist thinkers.

Section 2: Two Routes to Pluralism

In the previous chapter I mentioned the five principles of methodological pluralism identified by Sankey. These were:

- "Multiple rules: scientists utilise a variety of methodological rules in the evaluation of theories and in rational choice between alternative theories.
- Methodological variation: the methodological rules utilised by scientists undergo change and revision in the advance of science
- Conflict of rules: there may be conflict between different methodological rules in application to particular theories.
- Defeasibility: the methodological rules, taken individually rather than as a whole, are defeasible.
- Non-algorithmic rationality: rational choice between theories is not governed by an algorithmic decision procedure which selects a unique theory from among a pool of competing theories." (2000 p.213.)

To my mind, these principles are not all on a par with one another. I think of principles 1 and 4 as representing two fairly fundamental sources of plurality in scientific method, while 2, 3 and 5 can be seen as implications of 1 and 4.

When we say that there are multiple rules in science there are at least two sorts of things we might mean. On the one hand it might be the fairly trivial point that any attempt to explicate scientific methodology in detail will feature more than one of what Sankey and Nola call 'M-principles'. For example, we might have both 'prefer simpler theories' and 'prefer broader theories'. The more interesting reading of the claim that there are multiple rules in science however is the idea that there are multiple entire sets of methodological rules that each collectively make up the methodology of a field at a particular moment in the history of science. If that is correct, then the scientific method, in the general sense used in philosophy of science, of one field at one time may well be different from the scientific method of another field or another time. On this picture a full working out of the methods underlying theory choice in one area may well be incompatible with the methods underlying theory choice in another field or another time. Thus we would be faced with a pluralism of scientific methods.

The claim that methodological principles are defeasible is compatible with the general idea of there being only one correct set of general methodological principles in science. However, if these principles are defeasible, then there are situations in which they do not apply. This raises the question of how we are to know whether or not they apply in any given situation. It might be that this could be spelled out in advance by means of further principles governing the application of the first set of principles. However, according to Sankey and Nola, this is not what Feyerabend has in mind when he insists that methodological rules are context sensitive and not universalizable: "... in Feyerabend's view such conditions cannot be fully set out in some antecedent in advance of all possible applications of the hypothetical rule; at best such conditions are open-ended and never fully specifiable... even if rules appear to be universal... there will always be vagueness and imprecision concerning their application" (Sankey and Nola 2007 p.306).

A possible implication of the idea that methodological rules are highly context sensitive is that scientists in different contexts may effectively be using different methods, even if on the surface it looks like both are following the same abstract methodology. Once general rules have been applied in particular contexts, the relevant contextual features of the scientific inquiry may mean that individual rules have to be interpreted differently or not applied at all. Since different contexts may mean the negation or re-interpretation of different rules, this may mean that a scientist in one context is not using the same method as a scientist in another.

The importance of this distinction will emerge in my discussion of localism in the realism debate in chapter 6. These two sources of plurality in scientific methods match up to two competing attempts to reframe realism debates at a more local level.

Below I present different thinkers whose work supports the existence of one or the other of these sources of pluralism in scientific method. As we will see, in many cases these thinkers may appear to undermine the possibility of any historical argument over scientific realism on the basis of the reliability of scientific methods. Later in this thesis however I will attempt to reconcile the insights of methodological pluralism with a localised version of the historically focused scientific realism debate.

Section 2.1: Multiple Sets of Rules

Section 2.1.1: Kuhn

Kuhn (1962) argued that most scientific work takes place within what he called a 'paradigm'. Sankey and Nola summarise paradigms as "the source of an agenda of research problems, of exemplars used to solve problems, of experimental techniques and evaluative standards" (Sankey and Nola 2007 p.286). Paradigms also contain the theoretical underpinnings behind how scientists see and understand the world. In 'normal science' the paradigm remains unquestioned and its methods and assumptions are applied to solve the 'puzzles' that the paradigm frames as being most central to the field of enquiry.

Periods of normal science are occasionally punctuated according to Kuhn by periods of crisis and revolution. A crisis is prompted by an accumulation of problems that cannot be solved from within the existing paradigm, prompting scientists to question the paradigm and begin searching for alternatives. Eventually scientists settle on an alternative paradigm and again begin a period of normal science in which they apply the new paradigm to solve the new set of puzzles suggested by the new paradigm.

A paradigm shift comes with various kinds of change in scientific world view. For example, Kuhn argues that scientists undergo a perceptual shift in a paradigm change, such that how they perceive the world to be is fundamentally altered by working in a different paradigm. Most importantly for us, Kuhn argues that methods are tied to paradigms in such a way that a paradigm shift leads to 'methodological incommensurability' between paradigms. 'Incommensurable' means that there is no common methodology between paradigms against which apparently competing theories can be judged according to the same standards. Methodological incommensurability is one of several types of incommensurability found across paradigms according to Kuhn, with semantic incommensurability being another important type.

Because of the lack of shared standards across paradigms, Kuhn argues that it is not possible for the choice of a paradigm to be made entirely on logical or empirical grounds, and likens the experience of shifting from one paradigm to another to a religious conversion (Sankey and Nola 2007 p. 286).

This seems to suggest that all there is to say about scientific method is paradigm relative for Kuhn, meaning that there are presumably as many scientific methods as there are paradigms, but no overarching scientific method to perspective from which to assess the relative worth of these paradigm relative methods. This is supported by Kuhn's description of paradigms "as the source of the methods... for a scientific community" (Kuhn 1962).

What are the implications of Kuhn's view of scientific method for the scientific realism debate? First, and most obviously, Kuhn's account of science is a type of methodological pluralism, so a global understanding of the STI and PI is incompatible with it. Which methods are being used in science are dependent on which paradigm we are working in, and paradigms can undergo change and are specific to individual scientific fields.

What we can hope for though is that some more localised form of the scientific realism debate is compatible with some form of methodological pluralism. If there are multiple methodologies at play in science, perhaps we can evaluate the reliability of each of these methodologies individually by their historical track record. In this way the core of the scientific realism debate would be preserved without having to somehow resurrect methodological universalism.

However, the details of Kuhn's account are difficult to square with even this modified version of the historical scientific realism debate for two reasons. First, Kuhn seems to leave open a significant place for relativism or even irrationalism in his view of scientific method and has often been read as such (Sankey and Nola 2007 p.288). If this reading is at all correct, then there is no room left for questions like 'does this method reliably lead to true beliefs' to be asked about paradigm relative methods. Such questions could only be asked from within a paradigm and their answers would have no force outside of their own paradigms. Second, Kuhn's paradigms are too narrow in subject matter and time period to be amenable to historical assessment of the sort attempted by the PI as an assessment of the STI.

With regard to the first of these issues, Kuhn appears not to have held the strongly relativist views he is sometimes attributed. Sankey and Nola point out that in later work Kuhn did suggest some paradigm transcendent scientific values. For example he claims scientists in widely differing fields value theories if "they yield predictions (which should be accurate and quantitative rather than qualitative); they permit puzzle-formation and solution; they are simple; they are self-consistent;

they are plausible (i.e. are compatible with other theories currently deployed); they are socially useful” (ibid p.289). However, Kuhn still thought that each of these values needed to be given a ‘weighting’ or priority order so that conflicts could be resolvable. This weighting can be idiosyncratic to individual scientists, leaving an inescapably subjective element to scientific methodology, which again, does not sit well with the idea that scientific methods can be shown reliable or unreliable by their track record.

Turning away from the issue of relativism for now, perhaps the greater problem with squaring Kuhn’s views on scientific method with the historical realism debate is that for Kuhn, the more fundamental sorts of theory change in the history of science are often also changes of paradigm, for example the move from Ptolemaic to Copernican astronomy or from Newtonian to relativistic physics. If methods derive from paradigms and paradigm change leads to a change of methods, then it is very hard to see how methods could possibly be assessed historically in a way analogous to how the PI attempts to assess the STI. Theory changes like the ones mentioned above are key fodder for the PI. However, since the post-revolution paradigm does not share the methods of the pre-revolution paradigm there is little interest in the claim that the methods of the previous paradigm proved to be unreliable. The methods we are using in the new paradigm are distinct from these, and there is no reason to think that the new methods will be unreliable just because the old ones were.

This issue applies generally here. If we want to localise the historical realism debate to particular methods, then we need a way to individuate methods with sufficient generality for the same methods to be found across theory changes. Otherwise PI style arguments to undermine those methods are ruled out by how we have defined what a method is. The inevitable consequence if we accept that is that it is not possible to assess the reliability of scientific methods empirically.

Below I explore another version of multiple-systems-of-rules based pluralism to see if that provides a more favourable perspective from which to salvage the historical realism debate. This account is Hacking’s work on Styles of Scientific Reasoning.

Section 2.1.2: Hacking

Another form of methodological pluralism is Hacking’s (1982) account of styles of scientific reasoning. Hacking’s account of styles of scientific reasoning (SSRs) takes its inspiration from the historian of science Alistair Crombie’s work on styles of scientific thinking. Crombie argues for the existence of six distinct styles in science. These are:

- “a) the simple postulation established in the mathematical sciences,
- (b) the experimental exploration and measurement of more complex observable relations,

- (c) the hypothetical construction of analogical models,
- (d) the ordering of variety by comparison and taxonomy,
- (e) the statistical analysis of regularities of populations and the calculus of probabilities, and
- (f) the historical derivation of genetic development.” (Crombie 1981, p.284)

These styles of thinking are the product of the history of science and came into being through slow evolution of earlier ideas before crystallising at certain moments in the development of science.

Hacking claims that styles of reasoning set the conditions for sentences within those styles to be meaningful, what it is for claims to be rational, and what would determine the truth or falsity of a claim within that scientific style. Since SSRs play this role in setting the rules of a discourse, there is no perspective from which to ask whether a style itself is rational or irrational, justified or unjustified. What it is to be those things is determined by styles themselves. In this way SSRs are self-legitimizing according to Hacking. Once we adopt a style, that style can be used to justify itself.

Hacking at times compares his SSRs to Kant’s categories, but rather than being timeless and necessary, SSRs are historically contingent (Kusch 2010 p.162). Hacking sometimes labels SSRs as the ‘historical a priori’. What Hacking means here is that although SSRs are historically contingent, once established they become central to our way of understanding the world and are able to serve as a basis for objectivity. Styles eventually become “a rather timeless canon of objectivity” (Hacking 2002 [1992], p.187).

To give more content to the idea of what a style actually *is* Hacking suggests a number of necessary conditions for something to be an SSR. First, it must introduce new “objects, evidence, sentences (new ways of being a candidate for truth or falsehood), laws, or at any rate modalities, [and] possibilities” (ibid., p.188). Kusch (2010) illustrates this using Hacking’s work on statistics (style e above) (Hacking 1992a). New sentences introduced by this style include ‘The gross national product of Wurttemberg in 1817 was 76.3 million adjusted to 1820 crowns’. This had no truth value before the development of the statistical style. New classes include things like ‘those with a high risk of succumbing to a certain disease’.

Another condition for being an SSR is the presence of ‘self-stabilizing techniques’ (Kusch 2010 p.161; Hacking 1992). For example, in the laboratory style, a self-stabilizing technique is “the adjustment of auxiliary hypotheses to save a theory” (Kusch 2010 p.161).

Hacking is ambivalent about exactly how many SSRs there are in modern science (Kusch 2010 p.162). In addition to Crombie's initial six, Hacking adds the laboratory style and in places seems to be open to the idea that there a number of other styles present in science.

For our current purposes, if we accept Hacking's account of SSRs the revision of the STI and PI that suggests itself is to index these arguments to a particular SSR. The STI and the PI each need to be about a single method, or sufficiently uniform set of methods, in order to be workable. Hacking's work suggests that we might identify 'method' in this sense with an SSR. We might then ask whether a particular SSR has proved reliable over the period of science in which it has been in use.

Hacking's SSR account is obviously a variety of methodological pluralism since scientific methods are constituted by SSRs and Hacking claims that there are multiple of these.

This suggestion has a number of possible advantages, but also some serious challenges. Dealing first with the advantages, Hacking's styles are well suited to take the role of methods in a historical argument for two main reasons. First, they are broad enough to *have* a significant history and take in a wide number of theories in different areas of science. This means the problem posed by Kuhn's paradigm relative methods in the previous section does not apply. Theory change does not necessarily mean a change of style, so the PI is not ruled out by fiat if we ask whether SSR relative methods are vindicated by the history of science. Second, Hacking's account of SSRs suggests a manageable number of styles together with some suggestions about how to identify and delineate them.

However, there are problems with trying to use SSRs in this context. First, Hacking himself insists that SSRs are self-legitimizing, which seems to remove any need for them to undergo historical assessment. Moreover, Hacking's SSRs are so broad and basic that it is hard to imagine how we could subject them to historical scrutiny without using them to evaluate historical evidence, in which case our justification for them will be circular anyway. This situation would just serve to illustrate Hacking's contention that styles are self-legitimizing. Little would be added by a positive historical assessment of an SSR based on a method that used it, and a negative historical assessment of an SSR that relied on that SSR would be self-defeating.

Second, Kusch alleges that Hacking's position leads to a form of epistemic relativism. If he is right then it seems implausible that Hacking's SSR concept can be assimilated in any way with the realism debate, since both realists and anti-realists assume that whether belief in the truth of a theory is justified is not in any important way relative to anything.

Third, Kusch argues that Hacking does not really have any good way to individuate SSRs. Kusch attacks each of three criteria Hacking suggests for what makes something an SSR. Against the idea that an SSR is something that introduces new objects, evidence, sentences and modalities, Kusch argues that this would permit some dubious things to count as SSRs. Kusch gives the example of Alston's (1991) defence of mystical perception to support this (Kusch 2010 p.170). Kusch argues that Alston's theory of mystical perception introduces new objects (gods, angels etc.), new types of evidence (e.g. mystical perception as evidence for various religious claims), and new types of sentences with their own truth conditions (e.g. 'God told me to do his will'). Further, Kusch claims that Alston's views have their own self-stabilising techniques. Kusch takes Hacking's failure to rule this out as an SSR to be evidence that Hacking has failed to come up with any good way of identifying SSRs.

Kusch goes on to argue that there is a fundamental issue with how broadly Hacking draws his canonical SSRs in terms of timescales. Kusch introduces a distinction between 'longue durée' styles and 'court durée' styles, where a longue durée style "spans across millennia and many different forms of inquiry" and a court durée style is "specific to a narrowly defined period and field – say, 'the experimental style of German introspective psychology between 1890 and 1920'" (ibid p.170). Kusch argues that "the longer durée we are opting for, the less properties the resulting style will have, and the thinner its concept will be." He illustrates this point by imagining a style supposed to include "Hero's steam experiments in Alexandria ad 120, Buhler's introspective thought-psychological experiments of 1907, and a CERN particle experiment in 2005" (ibid). He claims that there is too much diversity in such a list for any style they all share to possibly have "very many features". But the consequence of this is that if we want to explain the methodological choices, or more broadly 'moves in the game of science' made by any particular scientist, appeal to this sort of longue durée style will be of little help.

Kusch pre-empts the objection that Hacking allows the possibility that styles evolve over time, which explains how they are able to persist. Kusch counters that the explanatory work would still have to be done by a court durée phase of the evolving longue durée style. This seems to leave little remaining room for the broader conception of style.

Kusch's argument seems closely related to one made by Otavio Bueno (2012). Bueno notes the diversity in scientific methods in different areas of science. For example, medicine uses randomised clinical trials irrelevant to particle physics, and particle physics uses particle accelerators irrelevant to ecology. "If there were a method that could be applied to areas as diverse as quantum mechanics and paediatrics, it would be unable to preserve any specific features of these areas. Such a method

would not contain any particular traits characteristic from such domains and would not be particularly relevant to these areas” (2012 p.663). For reasons along these lines, Bueno favours a narrower understanding of SSRs than that advocated by Hacking.

However, if Kusch and Bueno are right, and SSRs (or scientific methods generally) are better understood at a narrower level, then one of the advantages SSRs seemed to offer over Kuhn’s paradigms is lost. If we understand SSRs this narrowly, then there is not enough unity over a sufficiently long period of time for historical assessment of SSR based methods to get off the ground. No room is left for any kind of historical realism debate understood as an empirical argument over the track record of scientific methods.

Section 2.1.3: Discussion of Multiple-Systems-of-Rules Type Methodological Pluralism

In the work of both Kuhn and Hacking we see two different but related accounts of scientific method that both situate scientific method within a historically contingent framework. In both cases this framework is capable of changing over time and is not universal to science. Such views seem to imply that there is no one method found throughout science that could provide us with a unitary notion of success with which to justify the idea that lessons about science in general can be taken from the success or failure of the STI.

However, in both cases we face similar problems if we try to use the insights of Kuhn or Hacking as the basis for breaking up either the STI or the PI. Both philosophers have been accused of providing an account of scientific method that amounts to a form of epistemic relativism. It is hard to see how any kind of epistemic relativism could be compatible with scientific realism. Further, with both philosophers we face difficulty trying to individuate the frameworks that scientific methods are a part of in such a way as to make those methods amenable to historical assessment.

If we take either of these accounts seriously it is hard not to conclude that the whole project of assessing historically the reliability of a method of forming scientific beliefs is incompatible with methodological pluralism. I will return to this idea, and suggest a way around it after discussing a second source of motivation for methodological pluralism – the contextuality of methodological rules.

Section 2.2: Contextuality and Defeasibility of Rules

Section 2.2.1: Feyerabend

Using Galileo's arguments for Copernican cosmology as a case study, Feyerabend (1975) argued that there is no such thing as a universal scientific method. Feyerabend pointed out various aspects of Galileo's reasoning that appear to flout both the conception of scientific method that existed in his day and our modern conception of scientific method.

Feyerabend argues that Galileo made use of rhetorical devices, propaganda, and *ad hoc* arguments to defend his position. For example, Feyerabend argues that Galileo relied on *ad hoc* principles such as the principle of circular inertia in his arguments to explain why we don't observe the effects of the earth's movements, for example by observing a stone dropped from a tower falling away from the tower as the earth spins during the stone's fall. Feyerabend reconstructs the principle of circular inertia as follows: "an object that moves with a given angular velocity on a frictionless sphere around the centre of the earth will continue moving with the same angular velocity forever". (Feyerabend 1975.) This would explain why the stone falls at the base of the tower as the stone continues the circular momentum it had from the movement of the earth before being dropped, and therefore keeps up with the tower. However Feyerabend argues that Galileo's reasons for adopting it were largely because it is convenient to defend heliocentrism. This makes the principle *ad hoc*, a feature that many accounts of scientific method frown upon.

However, Feyerabend's conclusion was not that Galileo was suspect as a scientist. Rather Feyerabend argues that whatever methodological principle we might come up with in science, there will be circumstances where that principle ought not to be applied. Sankey and Nola summarise his position as follows:

"(I1b) Rules of method are highly sensitive to their context of use and outside these contexts have no application; each rule has a quite restricted domain of application and is defeasible, so that other rules (similarly restricted to their own domain and equally defeasible) will have to apply outside that rule's domain." (Sankey and Nola, 2000 p.35.)

This conclusion of Feyerabend's appears to present a considerable threat to the idea that 'success' has one meaning in science since if he is right that methodological rules have no application outside of a narrow domain then different areas of science should have different methods, and therefore different standards for evaluating the degree to which a theory is successful.

The interpretation of Feyerabend presented by Sankey and Nola differs from the common one in which Feyerabend appears to reject all ideas of scientific method by proposing 'anything goes' as the only defensible methodological principle. Sankey and Nola argue that 'anything goes' was meant as a position that a defender of universal scientific method is pushed into as that is the only principle that

could be viewed as universal. In a later defence of *Against Method*, Feyerabend clarifies that he does not advocate proceeding entirely without method in science, he just opposes regarding any methodological principles as universally applicable (Feyerabend 1978).

In relation to Galileo, Feyerabend appears to be suggesting that in Galileo's particular 'context of use' the methodological principles Galileo appears to have violated according to Feyerabend are not applicable. This would be an important finding as some of the principles Feyerabend accuses Galileo of violating appear quite fundamental. For example, Feyerabend accuses Galileo of using deliberately unsound arguments and rhetorical tricks to defend his position. If Feyerabend is willing to defend such egregious departures from the general image of good science however, it does leave one to wonder what exactly even highly context sensitive rules of method might be that would condone such behaviour. This makes it hard not to feel that, even if Feyerabend is on to something with the claim that methodological rules are highly context sensitive and defeasible, he may have slightly too lose an understanding of what those context sensitive rules might be.

Chalmers (1986) criticises Feyerabend's argument by arguing that Feyerabend has misrepresented Galileo's methods. Chalmers argues that Feyerabend emphasises Galileo's use of rhetoric and epistemological tricks too heavily partly as a result of misinterpreting Galileo's arguments and their context, and partly by focusing his attention on a work that Galileo wrote mostly in order to convince non-scientists of his position, i.e., *Dialogue Concerning the Two Chief World Systems* (1630).

Chalmers takes particular issue with Feyerabend's contention that Galileo helped himself to *ad hoc* hypotheses in order to defend Copernican astronomy. Chalmers argues that the principle of circular inertia, which Feyerabend claims is *ad hoc* can be seen to not be so when we look at *Two New Sciences* (1638). Chalmers argues that Galileo's adoption of this principle was not *ad hoc* because Galileo suggested ways to test it independently of the question of whether the earth moved around the sun, such as dropping an object from the mast of a moving ship and observing where it falls. Galileo's principle of circular inertia is also not *ad hoc* with regard to Copernican astronomy because the means by which Galileo arrived at the principle was unrelated to the debate over heliocentrism. Rather, Chalmers argues that this was not a principle at all in the sense of a fundamental part of Galileo's theory, but rather a consequence of his understanding of the laws governing falling bodies. Chalmers argues it is those that are central to Galileo's physics.

Chalmers' case that Feyerabend has unfavourably misrepresented aspects of Galileo's science is convincing. However, Chalmers goes on to argue that a proper understanding of Galileo actually lends even more support to Feyerabend's argument than the Galileo presented by Feyerabend.

Chalmers argues that Galileo's work demonstrated in practice that various methodological ideas from the time were incorrect: "Galileo's science was not guided by the methodological norms of its time and its success constituted a challenge to the adequacy of the standards they embodied." (Chalmers 1986 p.25) For example, Galileo's applications of mathematics showed the possibility of a mathematical physics which Aristotelian science had thought not possible, and the use Galileo made of the telescope was important in convincing others of its suitability for scientific observations. Most of all though, Chalmers stresses Galileo's introduction of experimentation as a major shift in scientific method, arguing that experimentation in its modern form was limited prior to Galileo. Chalmers concludes: "It is implausible... that such moves could be anticipated and legislated for in advance by methodologists. We can expect methodologies to alter in the light of new discoveries." (ibid)

Chalmers' characterisation of what can be learned from Galileo's defence of Copernican astronomy does not seem to fit with the idea that the defeasibility of rules is what leads to methodological pluralism. Chalmers seems instead to be noting that what the methodological rules are changes in the course of scientific development. In this way his position seems more an instance of multiple-rules type pluralism.

Section 2.2.2: Norton

Another philosopher whose work could be taken as providing support for the idea that methodological rules are highly context sensitive is Norton (2003). Norton gives what he calls a material theory of induction. He claims: "Contrary to formal theories of induction, I argue that there are no universal inductive inference schemas. The inductive inferences of science are grounded in matters of fact that hold only in particular domains, so that all inductive inference is local." (Norton 2003 p.647.) Norton argues that there is a tension in theories of induction between applying universally to all instances of induction, and properly describing how induction functions. Most formal theories of induction manage to find schemas that apply universally, but thereby fail to capture the force behind actual inductive inferences. Norton argues that whether any induction is valid or not is too intimately connected to the degree of warrant for our beliefs about the 'underwriting facts' for that induction. As an example of what he means he points to the following two formally identical inductive inferences:

"Some samples of the element bismuth melt at 271 [degrees Celsius]. Therefore, all samples of the element bismuth melt at 271 [degrees Celsius]" and "Some samples of wax melt at 91 [degrees Celsius]. Therefore, all samples of wax melt at 91 [degrees Celsius]" (ibid p.649).

The reason for the difference between these two inferences is down to the nature of our knowledge about 'bismuth' and 'wax' respectively. 'Wax' is a generic term for various substances while 'bismuth' names only one type of substance, which we are warranted in believing always has the same melting point.

Norton notes that formal schemas of induction need to fill in the conditions in which inductive inference can and cannot be applied. But Norton argues that there is no way to do this in such a way that the resulting formal schema remains an accurate account of *all* good inductive inferences. Instead, "the admissibility of an induction is ultimately traced back to a matter of fact, not to a universal schema" (ibid p.650). It is due to facts about bismuth and wax that there is a difference in the validity of the inductive inferences described above, and not to do with anything that can be captured by a universal schema of inductive inference.

Norton goes on to detail various attempts to "identify facts about the world that could underwrite induction" (ibid p.651), such as Mill's (1916 [1872]) principle of the uniformity of nature and Russell's (1912 chapter. 6) further working out of this. Norton claims that "All these efforts fall to the... irresolvable tension between universality and successful functioning" (ibid). Norton goes on to illustrate this point with regard to all of the major theories of induction found in the literature. He argues that "the more universal the scope of an inductive inference schema, the less its strength" (ibid p.662). In some cases, the basis for inductive inference can be so unique that they cannot be categorised with any other inferences.

Because Norton insists that inductive inference is so context sensitive, and the validity of an inductive inference depends crucially on our knowledge in the area where we are trying to use inductive inference, his work seems to tie in to the idea that the context sensitivity of scientific methods leads to methodological pluralism. Inductive inference is a crucial aspect of scientific method on any account of it. If the warrant for any inductive inference is as local as suggested by Norton, then the justification for inductive inferences varies enormously depending on which area of science we are talking about.

Further, the tension Norton points to between universality and adequate functioning in theories of induction seems to apply just as much in discussions of scientific method, as we have already seen several times in this chapter. This raises the prospect that his arguments about induction might generalise to scientific method in general, leading to the view that to properly capture the method of any area of science we will need to take a highly localised view of scientific methods. An implication of this view might also be that an intimate knowledge of particular subject matters is always necessary to assess the strength of the case for any given theory, and that abstract

arguments concerning the reliability of methods will not take us very far. As we will see in the next chapter, this kind of idea is taken up by philosophers such as Park (2019) and Fitzpatrick (2013) who argue for a highly particularist account of scientific confirmation that they explicitly use to undermine the historical scientific realism debate.

Section 2.2.3 Maddy and Fine

A philosopher who actually takes a position in some ways similar to the Norton-inspired view I described above is Maddy (2001). Maddy is concerned with whether we should take a 'one-level' or 'two-level' approach to claims about the world. The first level here is that of ordinary empirical claims of the sort made by scientists. The second-level is an epistemological level where philosophers try to judge from some kind of higher perspective outside of science what attitude to take towards claims made by scientists. Maddy argues that the scientific realism debate attempts to operate at this second level, but that the only legitimate level is the first level. Discussing Boyd's (1983) engagement with van Fraassen, Maddy says:

"... buying into van Fraassen's perspective tends to push Boyd away from the details of the local debate over atoms and towards global debates over such questions as whether or not the theoretical terms of mature scientific theories typically refer. The naturalist is wary of such blanket assertions, given the complexity of actual science: the particularity of arguments for the existence of individual theoretical entities, like atoms or quarks; the subtle gradations in levels of belief in the various parts of science; the widespread use of idealizations and mathematizations; and so on. At least at the outset, it seems unlikely that a single attitude towards 'the posits of mature science' will be correct across the board."
(Maddy 2001)

Maddy advocates engaging only with the local details of scientists' arguments and evidence for particular scientific claims and theories, and rejecting any attempt to argue in general terms for or against a general epistemic assessment of science as a whole. This view of which parts of science we should believe is shared by Fine (1986), who argues for what he calls the 'natural ontological attitude' (NOA). He says:

"All that NOA insists is that one's ontological attitude towards ... everything ... that might be collected in the scientific zoo (whether observable or not), be governed by the very same standards of evidence and inference that are employed by science itself." (Fine 1986)

The degree to which Fine and Maddy are motivated by the same sort of considerations I ascribe to the Norton inspired particularist about scientific methods is unclear, even if their conclusions look

similar. Both seem to approach the issue by thinking about the nature of truth and whether the proper standards for evaluating scientific claims are ever 'beyond' science itself. However, the quote above from Maddy does seem to mirror quite closely the sort of consideration I have been discussing, in which the diversity of methods, contexts, types of evidence and types of argument in science makes general assessment of the reliability of scientific methods difficult.

Section 2.2.4 Saatsi on Form and Content Driven Arguments

Saatsi (2009) draws a connection between Norton's theory of induction and the scientific realism debate. Saatsi distinguishes form-driven and content-driven arguments for realism. Form-driven arguments "appeal to the form of inductive inferences, whilst content-driven arguments appeal to their specific content" (Saatsi 2009). Saatsi identifies the No-Miracles argument as being "extremely form-driven" while the arguments scientists themselves use are "fully content-driven".

Saatsi uses reasoning inspired by Norton's material theory of induction to argue that content-driven arguments in the realism debate are to be preferred. Saatsi argues that form-driven arguments suffer from what he calls the "*Description – Justification Gap*" where "Too much emphasis is paid on formal descriptive unity, without realising that descriptive unity can be cheap and does not amount to justificatory unity" (ibid). The worry is that by focusing on the form of scientific inferences we miss out the material assumptions that are crucial to determining whether those inferences are justified or not. It also raises the risk that "an inductive inference is taken to be licit [in virtue of its form] when there is no relevant material fact to underwrite it" (ibid).

This focus on content over form seems to leave little room for historical inductions on the history of science. However Saatsi tries to leave some room for philosophical contributions to scientific debates when he says that "although the specific content-driven arguments hang on case-dependent detail, its master plan can be described in general terms" (ibid). The idea here seems to be that once we have given different content-driven, local arguments for different scientific claims, we can begin to compare these and identify patterns that can help us identify licit as opposed to illicit inductions. It can also help us to recognise whether inductions to the observable and some controversial induction to an unobservable are appropriately similar that we have to take the same epistemic attitude towards both. However, this kind of philosophical engagement with science seems to stop some way short of using the history of science to learn lessons about which sorts of methods and inferences in general are reliable. In order to preserve those arguments it would seem we need to have some way to defend the idea that arguments for realism or anti-realism that are to some extent form-driven can be workable.

Section 2.2.5 Stanford on the Particularist Threat to the 'Middle Path' Realism

Debate

Stanford (2021) highlights particularism about scientific confirmation as the main threat facing the shared project that participants in the scientific realism debate are engaged in. Particularists hold that first order scientific evidence always outweighs whatever evidence might be offered by the historically focused realism debate. The sorts of “generalities and abstractions” found in the scientific realism debate are “simply insensitive to precisely the sorts of variation in the details of the evidence we have in different cases that the particularist thinks really should generate varying degrees of confidence concerning various claims about the existence and character of fossils, or dark matter, or electrical charge, or bacterial infections” (Stanford 2021).

The position that my discussion of methodological pluralism has been pushing towards seems to be the one described by Stanford as particularist. Versions of pluralism inspired by the contextuality and defeasibility of rules in particular seem to lead towards the idea that little of use can be said about abstractly conceived methods and inferences.

Stanford seeks to respond to the particularist in two ways. First, he argues that what particularists recommend is what scientists are already doing, i.e., evaluating the first order evidence alone without too much concern for general principles of confirmation derived from the history of science. However, such scientists have often fallen into error, showing the potential usefulness of learning from the history of science in order to make our understanding of evidential relations better. Second, he points to examples of work that he thinks each represent “a clear example of epistemic guidance intermediate in generality between [the classical ambitions of traditional realists and instrumentalists] and the radical particularist’s competing conviction that useful epistemic guidance for science can only be found in the details of the evidence” (ibid).

The examples Stanford points to include some of his own claims about circumstances in which eliminative inference is likely to be more or less of a serious problem. Broadly, I think Stanford is moving in the right direction with the level of generality he aims at. I discuss his work in more detail in the next chapter and in chapter 7. Stanford’s phrasing of the issue strikes me as good however. What we need is we want to maintain the idea that substantive epistemic lessons can be learnt from the history of science is a way to resist the urge to highlight the variability, context-sensitivity and plurality of scientific methods to such a great extent that any abstract principles of scientific method, reasoning or confirmation become impossible to discuss or evaluate. The way I will be suggesting we do this will involve trying to identify methods and types of context in such a way as to allow historical assessment of methods relative to types of context.

Section 3: Lessons of Methodological Pluralism for the Scientific Realism Debate

To recap, I argued in the previous chapter that a global view of the scientific realism debate is dependent on universalism about scientific method. As we have seen, universalism about scientific method is not popular and for good reason. The tension Norton highlights between universality and proper functioning in theories of induction is a theme we have seen recurring throughout theories of scientific method as well, and this is taken to an extreme by universalism. Very little that is genuinely universal can be said about scientific method, and if anything can be said it is likely to be too vague to properly capture the underlying logic behind any particular science's methods.

My hope, which I will be continuing to explore in the remainder of this thesis, is that something important can be salvaged from the realism debate even if we move away from universalism and towards pluralism. My idea here is that we can find a version of methodological pluralism in which it is possible to run historical assessments of methods, similar to those done in the traditional realism debate, but on non-universal scientific methods.

However, my survey of some important methodological pluralist views above raises problems for this idea. In one way or another, each of the methodological pluralist views above seem to undermine the possibility of historical assessments of methods. Generally, this is because they either lead to a form of methodological relativism or because they conceive of methods as being so specific that they do not have wide enough applications or extensive enough histories to *have* a track record large enough to be judged against in the way the PI tries to judge the STI. These themes will come up again in the next chapter, in which I assess the current literature on localism in the scientific realism debate.

One thing that stands out from the discussion above is that if we want to salvage anything from the historical realism debate, we need to avoid particularism about confirmation. A view analogous to Norton's material theory of induction but referring to scientific methods in general leaves too little room for empirical assessment of methods. This is because it implies that scientific methods are so specific that it is impossible to make generalisations about the reliability of any abstract method. If that is the case then there is very little point in arguing about the reliability or otherwise of scientific methods considered abstractly. However, we need to be talking about scientific methods in the abstract in order for their track records to be enlightening with regard to their reliability.

However, multiple systems of rules type pluralism is also challenging to square with the historical realism debate. This is because it is hard to frame systems of rules in such a way that they persist through theory change. This is in part because there seems to always be pressure to conceive of

scientific methods in a narrower and narrower way, but some level of generality is needed to maintain any historically focused debate about the reliability of methods.

Turning aside from the challenges of framing a workable localised realism debate, it is worth noting that whatever challenges face localism in the scientific realism debate, we should not regard globalism as an option. Universalism about scientific method is not seen as a live option in contemporary philosophy, and rightly so. Science is too diverse for its methods to be captured by a unified theory, so we need to index any argument about the ability of scientific method to provide truths about the world to a particular method, and not expect these to lead to any general epistemic attitude towards science as a whole. One way or another, we *need* a good way to localise the realism debate that meets the challenges presented by the failure of universalism about scientific method. What this might look like is something I explore further in the following chapter.

Chapter 6: Localism and Individualism: Choosing the Right Level for Evaluating Historical Arguments in the Scientific Realism Debate

Introduction

In the previous chapter I argued that universalism about scientific method fails, and that therefore global approaches to the scientific realism debate are unsustainable. The diversity and context sensitivity of scientific methods suggests that we need to engage with the realism debate in a more local way. In this chapter I will review the work of some philosophers who have argued in favour of taking a more local approach to the scientific realism debate.

I will outline two strains in this literature – there are some who think that we should evaluate each theory entirely on its own merits and not expect historical evidence to be relevant to the epistemic status of current science. Park (2019) calls these thinkers ‘individualists’ (and argues in favour of individualism). Others think that it may still be possible to have arguments in the scientific realism debate that do call on historical evidence and offer conclusions for a part of science above the level of individual theories, but below the level of science as a whole. Park calls this position ‘localism’, or ‘collectivism’ (since it allows for theories to sometimes be evaluated collectively by way of historical inductions). Asay (2017) seems to hold this position.

I will attempt, eventually, to rule against individualism and in favour of allowing some historical evidence to be relevant in realism debates. First though, I present some arguments from Asay (2017), Park (2019), Fitzpatrick (2013) and Henderson (2018) to set up some central features of the localism issue. I will argue, following Henderson, that individualists lose some important resources for justifying instances of scientific reasoning, as they lose the ability to compare these to less contentious pieces of reasoning. I go on to argue that individualists may still be vulnerable to a historical anti-realist argument, namely Stanford’s (2006) new induction on the history of science. Individualists point to the particularity and context sensitivity of different inferences, pieces of reasoning, types of evidence and methods. However, even if these scientific arguments *are* highly context sensitive, the anti-realist may be able to argue that we have not always proven reliable at knowing whether the context we are in is one in which some methods apply. For example, Stanford argues that scientists have often tried to apply eliminative inferences in circumstances where this sort of reasoning could not have been reliable – they failed to recognise their context as being one where eliminative inference is unreliable.

I argue that to counter this sort of argument we need to be able to categorise scientific contexts in such a way that we can a) know what sort of context we are in and b) apply historical inductions to

find out whether a given method is reliable in that context. In the next chapter I will argue that such a categorisation can be achieved if we categorise scientific contexts according to the types of difficulty found in them.

A further lesson I draw from Stanford's argument about eliminative inference is that it is sometimes possible to identify in practice aspects of scientific method (in this case eliminative inference) in a way that seems well founded and avoids many of the difficulties with individuating methods I identified in the previous chapter. I argue that this may serve as one model of how to identify methods in a way that makes them amenable to historical assessment.

Section 1: Arguments for Localism

Section 1.1: Asay

In his paper "Going Local: A Defence of Methodological Localism about Scientific Realism", Asay outlines two ways of approaching the realism debate, which he calls 'globalism' and 'localism'. A claim is global if it applies "equally well across the board to all the various sciences", whereas a claim is local if it applies to some individual science rather than to all of science. A localist is one who thinks the appropriate level at which to conduct the realism debate is at the level of individual sciences, whereas a globalist is one who thinks the appropriate level at which to conduct the debate is the level of science as a whole. For example, one might be a localist scientific realist about chemistry but not biology. A globalist by contrast would think that we should either take a realist or anti-realist attitude towards science in general.

Asay believes that we should take a localist attitude towards the realism debate and presents a number of arguments for this conclusion. Park (2019) labels these:

1. The argument from diversity
2. The argument from disunity
3. The argument against hasty generalisations
4. The argument from other fields of philosophy

For the argument from diversity Asay points out that there is significant diversity across scientific theories in terms of what sorts of ontological commitment they involve accepting, and what sorts of epistemological issues are raised by those ontological commitments. This leads Asay into what Park labels the argument from other fields of philosophy. Here Asay notes a contrast between the ontological diversity in the objects of debate in the scientific realism debate and the apparent unity in the objects of other realism debates in other fields of philosophy. As examples of this Asay points to mathematical realism and moral realism. Asay claims that science is "nowhere near as unified a

domain of our thought as are morality and mathematics” (Asay 2017 p8). The cost of this is that it becomes harder for us to capture relevant nuances in the positions we take on the scientific realism debate, nuances which are more possible in realism debates that have more unified objects. Asay argues that a consequence of the diversity of science for the realism debate is that it creates a sufficiently large disanalogy with other realism debates that “we lose the ability to draw analogies with other forms of realism in philosophy” (ibid). Asay holds that this could be a useful tool in the scientific realism debate, and that by narrowing our focus to a particular field of science rather than science as a whole we put ourselves in a better position to exploit this tool.

The next argument presented by Asay is the argument from disunity. Here Asay draws on the work of philosophers such as Dupre who argue that science is disunified in the sense that different sciences do not reduce to one another, but rather that each has its own autonomy in terms of positing entities, developing methodologies and so on. When discussing methodological disunity, Asay very briefly seems to allude to an argument similar to the one I made in chapter 4, when he says, “If different sciences use different epistemological methods, then we need to evaluate those methods on a case-by-case basis in determining what the best attitude is to take with respect to them *vis-à-vis* realism” (ibid). He further argues that rejecting reductionism (where special sciences all eventually reduce to physics) shuts down “one possible avenue to globalism about science”, since realism about all the other sciences would piggyback on realism about physics.

Finally, Asay seeks to show by examples that in some cases philosophers have attempted to establish conclusions about science in general using arguments that are better suited to showing something only about some particular area of science. He uses Worrall’s discussion of theories of light as an example. Worrall analyses the shift from Fresnel’s theory of light to Maxwell’s and notes that while the entities posited by the two are different, there is something about the structure of the two theories and the equations they use that is preserved across the change. Worrall uses this to argue that we may be justified in being realists about the structure of successful theories even if we are not justified in being realists about the entities they posit. Asay objects to this however on the grounds that Worrall forms a conclusion about the whole of science on the basis of a single example. There are many other examples in the history of science that if allowed to form the basis of a global view would point towards different conclusions. The localist alternative suggested by Asay protects against these kinds of hasty generalisations by encouraging us to evaluate each case on its own merits: “If the history of the theory of light is best understood along the lines that Worrall tells it, then perhaps a structuralist view about *light* is appropriate; but it’s no argument at all that we should be structuralists about *genes* or epidemiology or any other areas of science” (ibid p.16).

Section 1.2: Park's objections to Asay

Park (2019) objects to Asay on the grounds that his arguments support a position Park labels 'individualism' rather than the localist view Asay aims to support. Rather than divide approaches to the scientific realism debate into globalist and localist, Park divides them into 'collectivist' and 'individualist'. A collectivist holds that it is possible to evaluate multiple scientific theories all together at the same time. Globalism is the most obvious example of a collectivist approach since the globalist thinks they can evaluate all scientific theories at once by way of arguments like the pessimistic induction. However, localism is also a collectivist position. The way Asay characterises localism, entire scientific disciplines are capable of being evaluated collectively. Since any scientific discipline involves multiple theories, this makes localism a form of collectivism. Individualism by contrast holds that the proper unit of evaluation in science is the level of individual theories. Park says that individualism in this context is committed to the view that "scientists' arguments for scientific theories are all the evidence for scientific theories, so the epistemic status of scientific theories is not affected by philosophical arguments, such as pessimistic... inductions" (Park 2019 p.3).

An important contrast between localism and individualism is that a role can be found for the PI and probably also for the STI by localism, but not by individualism. Localism allows for the formation of local pessimistic inductions. For example, say we are a local realist about biology. A PI could be formed stating that successful theories in biology have been false in the past, and that therefore success is not a reliable guide to truth in biology. Whether such a PI stands up or not, this is at least an allowable type of argument for localism as it is described by Asay. However, the whole point of individualism is that it denies any sort of argument that seeks to evaluate multiple scientific theories at once is allowable. No form of PI can be relevant for an individualist.

Park claims that each one of Asay's arguments in fact supports individualism and not localism. With regard to the argument from diversity, Park notes that "ontological and epistemic differences exist not only between present theories belonging to different disciplines but also between past and present theories belonging to the same discipline" (ibid p.5). As an example of this Park points to the caloric and kinetic theories of heat, which he claims differ just as much in their ontologies as theories from different disciplines might. If ontological diversity undermines global pessimistic inductions, why doesn't it undermine all pessimistic inductions?

With regard to the argument from disunity, Park argues that scientific change in a single discipline usually brings some degree of methodological change as well. For example, Park claims that Ptolemaic scientists and Copernican scientists used different methods, even though they were both

researching a similar set of questions and certainly fall under the same discipline. If methodological diversity across different fields of science undermines globalism, then methodological diversity in the same discipline across time should rule out localism in favour of individualism. This argument obviously echoes the difficulties I outlined in the previous chapter with using views like Kuhn's or Hacking's to frame local historical arguments in the realism debate. Methodologies that are relative to a local scientific framework change when those frameworks change, meaning that such methodologies generally have too short a lifespan for historical assessment of them to be possible.

As far as Asay's argument from hasty generalisations goes, we run the risk of hasty generalisations according to Park with any kind of collective evaluations of theories: "the fallacy of hasty generalisations occurs not only when our generalisation is about all theories across different disciplines but also when it is about all successive theories in the same discipline." (ibid p.6.) For example, Park finds the following generalisation too quick: "since theories of light were unstable in the nineteenth century, all theories of light will be unstable." Park does not defend this claim in great depth, but the idea seems to be that the spirit of wanting to avoid hasty generalisations should lead us to evaluate all theories on a case by case basis.

Asay does to a degree pre-empt Park's objections when he considers the idea that localism collapses into what Asay calls 'hyperlocalism', or "realism debates about the most narrow domains of science possible. In the limit, hyperlocalism reduces scientific realism to indefinitely many forms of realism about single, individual theses." (Asay 2017 p.19). Hyperlocalism seems to be much the same thing as Park's individualism.

Asay's response to this idea is to say that the localist, "need not have any principled problem with hyperlocalism", but that, "the domains to which we apply 'realism' and 'anti-realism' should be as broad or as narrow as is useful to us" (ibid). The important point for Asay is that treating science all the same draws the domain too broadly to be useful. If in some context it makes sense to evaluate scientific theories individually rather than collectively, Asay is fine with that.

However, Park's claims are too strong to be so easily reconciled. Asay still leaves open the possibility that in some cases scientific theories *can* be evaluated collectively, where Park wishes to establish that this is never permissible. Perhaps Asay's main intention is simply to refute globalism and he would be happy to endorse Park's individualism in favour of the localist position he sketches. However it seems worth pursuing further the question of whether localism is a viable option, particularly as we have seen that individualism is incompatible with any kind of historical argument in the realism debate. If we want to save anything of the existing historical debate over scientific realism, we need a way to make collectivism work.

Park's main positive argument against collectivism in general is that it clashes with what he calls a "basic rule of how to evaluate an argument" (Park 2019 p.8). His explanation of this rule is as follows: "Suppose that there are two arguments. One is intended to justify p ; the other is intended to justify q . If we want to decide whether to accept or reject q , we should evaluate the argument for q , and we should not evaluate the argument for p " (ibid). Park thinks that philosophers who use historical methods to engage in the scientific realism debate break this rule. Suppose that we are wondering whether we ought to be realists about some modern scientific theory. The arguments in favour of this theory are the arguments used by scientists, i.e., the first-order scientific evidence that led to the theory becoming widely accepted. They are represented by q above. For Park, any historical consideration concerning other theories, the arguments for them, or their eventual truth or falsity, are represented by p above, since they are not *directly* concerned with the truth or falsity of that theory.

Park's argument here does not seem to work. The relevance of the history of science in shaping our attitudes to current science is not hard to construct in a way that makes for a good argument. Park characterises the pessimist's method as being to argue that the arguments for a current theory's predecessors were bad, so arguments for the current theory must also be bad. Reconstructed thus, the pessimist's argument does seem suspicious. However, once we add the premise that the arguments for a theory's predecessors were generally similar to the arguments for the current theory, the argument is significantly strengthened. The pessimists' argument is really that past theories were formed using the same method as current theories, and that since past theories turned out to be false, that method has been shown to be unreliable. Pressure can be placed on this argument by disputing that the methods in question really are the same in both cases, but the argument is not of an impermissible type. As I explained earlier, we use these kinds of arguments to evaluate the reliability of a method in many areas of epistemology, not just in the scientific realism debate.

Of course, Park's argument makes more sense if we assume particularism about scientific confirmation as a premise. If we assume that scientific methods are too context sensitive to be reliable or unreliable in the abstract then it makes sense to think the success or failure of past theories irrelevant to the strength of the case for current theories. However, this kind of particularism is essentially the conclusion that Park is trying to establish. His argument therefore seems not to be capable of convincing anyone who does not already agree with him.

So much for Park's independent argument against collectivism then, but is he right that the motivations for localism *inevitably* lead us to go further and endorse individualism. In his objections

and replies section, Park puts his finger on an important point. He notes, correctly, that (collectivist) localism is happy for the proper unit of evaluation to vary in the realism debate. Sometimes the proper unit may be a whole scientific discipline, sometimes a sub-discipline, maybe in some cases a single theory, and so on. Depending on how we specify the unit of evaluation, pessimistic inductions may be allowable or not, and they may be supportable or not. However, little detail has been provided in the work of philosophers such as Asay about how the proper unit of evaluation is to be determined. As we saw in the previous chapter, specifying the proper level of evaluation in terms of scientific methods is not easy.

I agree with Park that there is a hole here in the case that exists for localism in the current literature that needs to be filled in. Park avoids the issue by keeping the unit of evaluation fixed – we should always evaluate scientific theories individually. Globalists also avoid the issue by keeping the unit of evaluation fixed – we should always evaluate science as a whole. However, (collectivist) localism's failure so far to fix the unit of evaluation leaves a gap. We cannot decide arbitrarily how to draw the boundaries when deciding how far to cast our realist net. We need some principled reason for including some parts of science in our discussion and excluding others. That some areas tend to fall under the same broad discipline does not seem like enough here. Imagine for example we decide that we want to be local realists about biology. How do we forestall pessimistic induction cases from physics? Asay's answer seems to be to point out the differences between physics and biology, but we need some further story about why *these* differences are important while the differences found *within* biology are not. Filling in this gap is a task I will attempt in the second half of this chapter and the next chapter.

Section 1.2: Fitzpatrick

Fitzpatrick (2013) argues that realists should abandon the STI in favour of what Fitzpatrick calls the 'local strategy'.

Fitzpatrick presents two main arguments in favour of the local strategy and against the need for a general no miracles argument. First, he argues that the local strategy is better placed to withstand the PI than a form of realism that relies on a global no miracles argument. Second, he argues that the role that the NMA is usually thought to have played in supporting realism as a framework is unnecessary.

The core of Fitzpatrick's local strategy is an insistence that the specific details of the evidence favouring different scientific theories matter when we are deciding whether the inferential methods used in forming those theories are reliable. Where the NMA and the PI both aim to establish

conclusions about the reliability of scientific inference in the most abstract possible terms, Fitzpatrick claims that only by attending to specific details of scientific cases that we can tell whether the inferences employed by scientists are likely to be reliable. Fitzpatrick cites Achinstein's (2002) discussion of Perrin's arguments for the atomic theory. Achinstein argues that the details of Perrin's case for the atomic theory renders standard anti-realist arguments not applicable.

General support for this line of thinking is given by Norton's material theory of induction, in which "the inductive inferences of science are not licensed by universal schemas. They are grounded in matters of fact that hold only in particular domains, so that *all inductive inference is local.*" (Norton 2003)

Fitzpatrick also points to the work of Psillos as a source of support for the local strategy as applied to inference to the best explanation (IBE). Although Psillos argues that 'the realist framework' needs justification as a whole, he also claims that "IBE-type reasoning has a fine structure that is shaped, by and large, by the context" (Psillos 2007).

If we accept both positions, and think that both inductive and abductive reasoning cannot be captured in abstract terms alone and without regard to contextual features of an epistemic situation, then the need for any global defence of realism appears to be greatly weakened. Since it is only by examining the details of the scientific evidence that we can know how strong the case is for realism about any given field, there is nothing to be gained from trying to justify realism as a whole. Thus Fitzpatrick's claim that the job done by the NMA, viz. justifying the 'realist framework' does not really need doing.

Fitzpatrick's second point, and, one suspects, his chief motivation in endorsing localism is that the localist has an easier time dealing with historical challenges to realism such as the pessimistic induction. By insisting that confirmation is as localised as Fitzpatrick says, and therefore by insisting that the methods of science are particular to the case for each individual theory, Fitzpatrick breaks the analogy between current theories and historical successful-but-false theories.

Using Park's distinction between localism and individualism, it becomes clear that the position that Fitzpatrick is pushing towards is really individualism, even though he calls it localism. Fitzpatrick's position is rooted in particularism about confirmation – in particular the idea that induction and abduction are both incredibly context sensitive. The degree of context sensitivity indicated by Fitzpatrick appears to align well with individualism, but does not necessarily leave open the possibility that a scientific discipline as a whole might be the appropriate level for realism debates.

His main example, that of Perrin, concerns the case for a specific theory rather than the methods of an entire field.

Fitzpatrick's arguments are inspired by the same kind of motivation as those of Feyerabend and Norton that I discussed in chapter 5. Fitzpatrick is right that if scientific methods are context sensitive to a very high degree this threatens the possibility of learning much about the epistemic status of current theories from the successes or failures of past theories, as if methods are context sensitive in this way this suggests that it is not possible to establish the reliability or unreliability of methods in the abstract.

Section 1.3: Henderson's Objections to Localism

Henderson (2018) is unconvinced by Fitzpatrick's attempts to avoid the pessimistic induction. She argues that the supposed benefit of avoiding historical challenges to realism is offset by the inability to provide a justification for scientific reasoning.

Henderson's argument here relies on the idea that the way to justify some particular piece of reasoning is to show that it is an example of a generally respectable type of reasoning. However, if confirmation is as particular as claimed by Fitzpatrick then this is not possible as little insight is gained by assigning some particular piece of reasoning to a general type.

Henderson argues that this means that rather than being in a better position with regard to historical challenges to realism, the localist is left "with fewer resources than the globalist for dealing with the historical anti-realist". (p.161) The globalist can hope to show that a particular piece of contested scientific reasoning has a "good track record historically". So if we are unsure whether the case for some particular theory is strong enough to justify realism, we can look at the historical pedigree of similar pieces of scientific reasoning. If this sort of reasoning generally has not led us astray in the past, this should strengthen our conviction that it is not doing so now. Without this kind of argumentative strategy to provide objective support for a piece of reasoning, it is hard to see how realism is to be justified about the products of that reasoning.

Henderson claims that this puts the localist in a difficult position. In order to give philosophical backing to a scientific theory they need to show that the methods used to form that theory were reliable. But this can only be done by identifying that the methods in question are particular instances of some generally reliable method. Once that has been done though, it is hard to avoid the question of whether the method in question has also been reliable in other instances. This opens the localist back up to the pessimistic induction

From these objections it seems that the target Henderson has in mind is what Park would call individualism, rather than a collectivist version of localism. It is also easy to read individualism rather than localism as the position that Fitzpatrick is arguing for. This is important, as to avoid Henderson's concerns we do not need to be globalists, we just need to not be individualists. If we believe that scientific methods are so context sensitive that each theory is essentially justified by a different method, then Henderson is right to say that we lose important resources for assessing the reliability of a method. We cannot look at the track record of any method to establish its reliability, as no method has been used sufficiently widely to have a track record. However, if we instead believe that scientific methods are often used across *some* variety of scientific contexts but without being universal, then we can still look at the track record of these methods. The important thing is to be sure that all the cases we are looking at really are instances of the same methods, and further that the method we are considering is applied only to relevantly similar topics. Narrowing the scope from 'scientific method' to 'the method found in these areas of science' and from 'science' to 'these relevantly similar scientific topics' does not necessarily mean embracing particularism about scientific confirmation.

However, even if Henderson's arguments do not necessarily undermine localism, they are good arguments against individualism. Drawing similarities between a piece of contentious reasoning and an agreed example of good reasoning is a crucial argumentative tool when it comes to justifying any piece of reasoning. Individualism seems to voluntarily give up this tool by rejecting the possibility of such comparisons being instructive. This seems like one good reason to prefer 'collectivist' localism to individualism as long as we can find a workable version of collectivist localism that avoids the problems I spelled out in this and the last chapter. In the next section I consider Stanford's Problem of Unconceived Alternatives as on the one hand a practical example of how to individuate a scientific method so as to make it amenable to historical assessment, and on the other hand as an illustration of the shortcomings in the individualist's attempts to resist historical anti-realist arguments.

Section 2: How the historically motivated anti-realist can still attack localism

In this chapter and the previous one I have presented several variations on the general idea that scientific methods are too context sensitive for it to be possible to evaluate their reliability in the abstract by assessing their track record. Philosophers such as Fitzpatrick and Park have claimed that such a position avoids the need to engage with historical anti-realist arguments such as the pessimistic induction. These historical arguments miss the point because they miss the fact that whether any kind of scientific reasoning is likely to be reliable is in large part down to the particular details of the scientific argument, rather than down to which formal schema for an inference type is

in play. It is what Norton would call the material postulates at play in scientific inferences that determine whether those inferences have been well applied and are therefore reliable. If the material postulate implicit in an inference is well justified then the inference itself will be equally well justified. However, material postulates are specific to their own theories, and it is very hard to capture any general principles for when they do or do not make an inference well justified. The applicability of a scientific method is not something that can be given by an algorithm or a universal set of conditions on such a view. So whether those methods provide a good basis for knowledge must be evaluated on a case by case basis and there is little to be learnt by assessing the track record of a method.

I do not think that this position really does avoid historical anti-realist arguments successfully. The core idea of the PI is the thought that in the past scientists thought they were in a position to infer the truth of their theories, and they were wrong – and that what it felt like to be those scientists is the same as what it feels like to be a current scientist. Past scientists thought that *their* material postulates were well justified and true and suitable to form a base for inductive inference etc. but they were mistaken. If so many past scientists were thus mistaken, why not think that current scientists are also mistaken about the solidity of their material postulates?

There is a very well-known recent argument in the scientific realism debate that illustrates what I mean here. This is Stanford's Problem of Unconceived Alternatives (PUA) (Stanford 2006). Stanford identifies a type of inference which he believes, rightly it seems, to be widespread in science. This is eliminative inference. Eliminative inference has the following form:

- One, and only one, of A, B or C must be true
- A is not true
- B is not true
- Therefore C must be true

Stanford argues that the history of science demonstrates that this type of inference is unreliable for a large part of science, as scientists have repeatedly shown themselves incapable of conceiving all relevant alternative hypotheses. That is, when framing the premise 'One, and only one, of A, B or C must be true', they have failed to think of alternatives D, E and F which are also genuine contenders for being the truth. Often, one of D, E and F turned out to be better supported by the evidence available at the time, and ended up superseding C. Stanford argues on this basis that we should not be confident we are not ourselves in this position with many current scientific theories. We have no way to be confident that we have not failed to conceive relevant alternatives to our current theories that are at least as well supported even by the evidence currently available to us.

I will leave the details of Stanford's argument, and discussion of how far the argument works, until later. However there are several general points about Stanford's argument that seem instructive. First, Stanford has demonstrated in practice how we can isolate a method within science such that it appears to be capable of historical assessment. This shows a way forward from the difficulties of individuating methods in order to assess them historically which I outlined in the previous chapter. Second, Stanford's argument still has force even if we do think that eliminative inference is highly context sensitive. This demonstrates the shortcomings in the particularist case against historical anti-realist arguments. I will argue that there *is* a good response to Stanford's argument, but that it is not compatible with individualism. My argument will support Henderson's contention that in order to defeat historical anti-realist arguments we need some degree of what Park calls 'collectivism'.

Regarding the first of these points, i.e., that Stanford has shown in practice a way to identify methods in such a way as to make them amenable to historical assessment, eliminative inference has two key features. First, it is broad enough to be used in a wide range of scientific contexts and across a significant period of time. This allows it to *have* a track record broad enough for us to have something to assess eliminative inference against. Second, eliminative inference is not relative to any specific scientific framework, in the same way as an overall methodology may be. In chapter 5 I outlined the accounts of Kuhn and Hacking as examples of methodological pluralism based on the existence of multiple systems of rules. From these accounts it can seem as if everything there is to say about a methodology is relative to a historically contingent framework, however, both Kuhn and Hacking allow for some inferences and other methodological principles standing outside of their historically contingent frameworks. For example, in his later work Kuhn argued for the existence of paradigm transcendent values in science (Sankey and Nola 2007). Similarly, Hacking does not see basic inductive or deductive rules of inference as style dependent or as constituting their own styles (Hacking 1982).

This suggests that there are at least some features of scientific method that are not relative to any specific scientific framework, but can be seen as methodological resources that are shared quite widely throughout science. Eliminative inference is such a resource. It is not dependent on any particular framework, but is a general form of inference found in many areas of both science and non-science. It is a methodological tool that is shared by many more specific scientific methodologies. That seems like a template for what we should be looking for when it comes to historical assessments of methods – methodological resources shared by many specific scientific methodologies, but without any need for these to be universal across science or part of some general, universal scientific method.

The status of eliminative inference as a shared methodological resource across many areas of science explains why there does not seem to be pressure to narrow down the method discussed by Stanford in the same way as there is pressure to narrow down methods in the sense of complete sets of rules for an area of inquiry, as discussed in chapter 5. Breadth is not an issue because Stanford is not attempting to give a complete account of any scientific field's methods, he is just highlighting one way in which different fields make use of a common argument form. Highlighting one dimension in which different scientific methods are similar and providing an argument that methods similar in that way will face some particular problem looks like a good way to go about historical assessments. Of course there will still be difference between the methods used in different areas of science, but if they all use eliminative inference then perhaps it is reasonable to think they all run into the characteristic difficulties of using eliminative inference.

Not trying to run a historical assessment of a complete methodology also avoids the confirmatory holism issue I talked about in chapter 4. There I worried that if we try to find out whether a method has generally led to true theories we will find that the historical cases we are considering differ in at least some ways from each other, so could not be counted as instances of the same method. I also worried that even if we did come to the conclusion that method we were considering was not reliable we would be faced with a vast choice of possible corrections and little idea of how to decide which one to go for. However, none of these worries apply if we are only interested in any methods that are similar to one another in a single specified way, or putting it another way, in any methods that crucially depend on the same methodological resource we are discussing.

How many methods in the sense of shared methodological resources are there for us to evaluate through an analysis of the history of science? Will this way of conceiving methods confine the realism debate to being about just a handful of different, fairly broad, inference types? I do not think so. The important thing when framing an anti-realist argument is I think to make explicit the importance of whatever methodological feature we are talking about. It seems to be possible in principle for any methodological resource to provide the core of an anti-realist induction on the history of science. The important points are just that, a) the methodological resource is found widely enough to have a significant historical track record, and b) some plausible argument can be made that there are special epistemic considerations raised by relying on that methodological resource in some fairly broadly construed scientific contexts.

Returning to the point I mentioned earlier that Stanford's argument bypasses objections based on the context-sensitivity of scientific methods, Stanford's argument is not about whether eliminative inference is a reliable method of forming beliefs in the abstract. It is obvious that eliminative

inference is reliable in the abstract, as eliminative inference of the sort I sketched above is *deductively valid*. His argument is that we have proved unreliable in knowing how to properly apply it in science because we have frequently proved incapable of conceiving all relevant alternatives.

Contextual factors are obviously important to whether an instance of eliminative inference represents a good argument. For example, local context sometimes seems to be able to narrow down the possible range of alternative hypotheses. In chapter 3 I discussed the choice biologists in the late nineteenth century faced between evolution and creation as accounts of the origin of species. There it seems as though we have good reason to think that there are only two options in the logical space. Either modern species were created in their current forms, or they are the result of a process of change acting on different earlier forms. Phrased at this level of abstraction, it is hard to see what other possibility there might be. However, this seems to be specific to this question and not easy to generalise to other scientific contexts. Achinstein's analysis of Perrin (Achinstein 2002) provides another example of how local context can help to constrain the range of hypotheses.

However, points like this only go so far in providing a response to Stanford's new induction. This is because the crux of Stanford's argument is that we are not in general good judges of whether we are in a context where eliminative inference is reliable or not. Or to put it another way, we are poor judges of whether we have exhausted the logical space. Past scientists often thought they had according to Stanford, only to later be proved wrong. So why should we think that we will not also be proved wrong if we think we have exhausted all alternative explanations in current scientific inquiries.

This challenge cannot be met by simply arguing that scientific methods and inferences are highly context sensitive, and that we cannot generalise about whether any given method *is* reliable or not. The question is whether or not we are able to recognise whether it is reliable or not in whatever context we might be in. What is needed to meet Stanford's argument is some analysis of the types of context in which eliminative inference is reliable, or rather in which we are reliable at applying it, and the types of context in which it is not. However, once we get into categorising contexts in this way, we open up the way again for historical assessment of a method *within* those contexts and are therefore back at a form of collectivism. In this way Stanford's PUA seems to create pressure away from individualism and towards some form of collectivism.

The way Stanford argues against eliminative inference shows a model for the historical anti-realist to argue against any inference or method in science in a way that nullifies the individualist's defensive resources. Even if we admit that scientific inferences and wider methods are highly context sensitive, it may still be that we have proven to be unreliable judges of what sort of context we are

in, and therefore of whether a given inference or method is reliable. If that is the case, it amounts to much the same thing as that inference or method itself being unreliable.

In the next chapter I look at how we might categorise contexts in order to frame the contexts within which we can generalise about the reliability or otherwise of a scientific method. I argue that scientific contexts are most usefully and plausibly categorised according to which sorts of difficulties we face in different areas of science.

Chapter 7: Individuating Topics in Historical Assessments of Methods

Introduction

In the previous chapter I argued that particularists about confirmation are not as well placed as they claim to be with regard to historical anti-realist arguments. This is first because, as Henderson argues, the particularist loses a key strategy for defending any contested piece of scientific reasoning in that they cannot draw comparisons between these and less controversial inferences. The second reason for this is that it is hard to see how the particularist can respond to any argument claiming that we have proved ourselves unreliable at knowing when to apply different sorts of inference or method. Stanford's PUA for eliminative reasoning presents this type of argument, and resorting to particularism about scientific confirmation does not seem to help matters a great deal. Further it seems like it may be possible to reframe other PI type cases so they have the same general form as Stanford's PUA and are equally difficult for the particularist to respond to.

A more promising way to meet historical anti-realist arguments is to attempt to use the importance of context to narrow their scope, but without claiming it is impossible to generalise at all about the reliability of scientific methods. What we need is to identify a way of grouping the different ways in which the specific context plays a role in the reliability or applicability of methods in order to allow us to generalise within these types of context. If we can identify a range of different types of situation in a way where we can easily recognise which type of situation we are in and generalise about the reliability of a method for that type of context then we should hopefully have the tools to resist anti-realist arguments in at least some areas of science.

In this chapter I aim to show how it is possible to group together different contexts in science in order to 1) limit the scope of anti-realist arguments such as Stanford's in a principled way and 2) show that this makes it possible for the history of science to provide evidence of the reliability or otherwise of scientific methods for particular contexts.

I begin by asking how the sceptical consequences of historical anti-realist arguments can be confined to whichever target a given anti-realist has in mind. I present several different characterisations of the target of anti-realist arguments, and point out that all are variations on the general idea that the answers to some scientific questions may be unknowable because they are too difficult. Exactly what sort of difficulty is held to be important varies however. I attempt to build on existing work on this issue by attempting a catalogue of different sorts of difficulty in science – I argue that each type of difficulty I describe is largely independent of the others, and it is a mistake to group them together in

terms of their likely epistemic consequences. There is no reason to suspect that any given method will be equally reliable or unreliable with respect to more than one type of difficulty.

I then move on to consider how widely we can generalise about the reliability of methods relative to a particular type of difficulty in science. I argue that if we relativise judgments about the reliability of methods to a type of scientific difficulty this undercuts some of the motivation for particularist views of scientific confirmation. Relativising to a type of difficulty does this by finding a role for the importance of specific contextual detail, without thereby ruling out generalisation across different specific contexts. Once we realise that in any actual scientific context there are likely to be multiple methods and multiple sorts of difficulty in play, we gain a lot of scope for each scientific theory to require some degree of independent evaluation and at the same time we still have the potential to learn about the reliability of methods from the history of science. This is particularly true as fully appreciating all methods and types of difficulty at play in a real scientific context will require substantial local knowledge. These lessons about the reliability of methods for particular topics in turn allows for general philosophical backing (or counter-arguments) to be given to scientific theories and claims. We have hopefully at least some scope to argue through the history of science that the methodological tools used to form and support a theory have generally been either reliable or unreliable for relevantly similar topics (i.e., topics showing similar types of difficulty).

Section 1: How can the Anti-Realist Characterise the Limits of their Scepticism?

Historical arguments for anti-realism about science generally have the character of pushing for the conclusion that it is impossible to know the answers to at least some scientific questions, as the history of science has shown that methods we must rely on to answer those questions are unreliable. For example, Stanford argues that eliminative inference has proven unreliable for quite a broad range of scientific questions. The traditional PI argues that inferring truth from success in science is unreliable. If they work, these arguments do not point out any failing in the epistemic position we are in *now* with regard to these methods. There is no future point we might get to where the success was so great that we could infer truth for example. Instead it looks like we could never be in a position to acquire knowledge through the method targeted by a historical anti-realist argument.

However, no form of anti-realism about science is intended to be a special case of general Cartesian scepticism. Scepticism is supposed to be confined to only a particular area, which calls for some account of how the anti-realist's sceptical arguments are to be confined to the intended area.

Reflecting on how this is generally done seems a good starting point in an attempt to separate out

different areas of science within which we can generalise about the reliability of a method, but between which generalisation is not possible.

Let's start with a rather naïve characterisation of anti-realism about science. According to this view, we are not in a position to know that the claims of any scientific theory are true. We should be sceptics about all of science, but about science specifically as opposed to non-science. It is not likely that many current philosophers hold any such view (see Vickers 2022 chapter 2 or Stanford 2021) – most current anti-realists are willing to accept that we are justified in believing at least some scientific claims. However it still seems illuminating to consider how such an anti-realist might confine the scope of their arguments, as I will explain.

The first challenge for this naïve or radical anti-realism is to reconcile their position with the diversity of science. It would be very surprising if the epistemic position we are in with regard to one bit of science is generally similar to the epistemic position we are in with regard to every other bit of science. The second challenge is to explain why the position we are in with regard to all the different areas of science is *worse* than the situation we are in with regard to large areas of non-science (since the anti-realist is not supposed to be a Cartesian sceptic). An explanation would be desirable both for why we are in such a bad position with regard to science as opposed to non-science, and also why we are in an equally bad position with regard to all the many areas of science.

What could supply such an explanation? It seems fairly clear that nothing to do with the methods of science compared to non-science will provide an answer here. First, methodological diversity across science is too great for methods to provide a means of unifying science compared to non-science (see chapter 5). Second, there is no obvious reason to think the methods of science are likely to be *less* secure than more everyday methods of forming beliefs.

The only thing I can think of to say here is that scientific questions are generally more *difficult* than, for example, everyday empirical questions like 'what time is the train to Durham?'. This extra difficulty could perhaps explain what scientific questions have in common in such a way as to make it plausible that no method could provide knowledge in science.

This seems like a step in the right direction as it at least gives us some reason to think that there will be something in common across scientific subject matters with regard to the likely reliability of a given method of forming beliefs – since scientific questions are so difficult maybe no method can meet the evidential standard needed for realist commitment. In effect it requires us to say that the answers to some questions are unknowable (regardless what method we use) because they are so hard. That view seems to be a central part of the anti-realist's position anyway.

Of course, to say that we cannot know the answers to scientific questions because those questions are hard is only a vague point in the direction of an explanation, rather than an explanation as such. At this level of generality, the explanation lacks much initial plausibility. Many questions in mathematics are hard, but doable. Finding a way to get from London to New York in less than 8 hours was hard before suitable aeroplanes had been developed. We now know how to do this. However I have chosen to start at this level of generality because it seems like it must be the crux of any explanation of the unknowability of scientific matters. It requires working out in some more specific form to be plausible, but any workable explanation will be *some* form of this general idea.

So what are the sources of difficulty in answering scientific questions? Well, there are many, and most are found only in answering some scientific questions, not all. This suggests that as well as narrowing our focus in the realism debate to particular methods, we also need to narrow our focus according to the type of difficulties we face in different areas of science. Just because a method is unable to give knowledge about questions showing some of these kinds of difficulty is no reason to suppose it will also be unable to give knowledge where only other kinds of difficulty are present.

Below I attempt a very brief catalogue of sources of difficulty in science. I deal with five main sources of difficulty:

1. Inaccessibility of the subject matter
2. Unfamiliarity of the subject matter
3. Conceptual complexity in our theories
4. The propensity of some subject matters to change over time, sometimes in response to scientific investigation.
5. Difficulties arising from form of question being asked

Section 2: Types of Difficulty in Science

Section 2.1: Inaccessibility

In *Exceeding Our Grasp* the main word Stanford uses to characterise the areas of science he thinks we should be sceptical about is 'inaccessible' (for e.g. 2006 p.158). What he is getting at is that some areas of science deal with subject matters that, for one reason or another, we cannot observe. Reflecting on this general characterisation, there are various reasons why something might be inaccessible for observation. Sometimes we cannot observe our subject matter at all, sometimes we can only observe a part of it, and sometimes we can only observe it with the help of machines like microscopes. The reasons why a subject matter might be in some way unobservable are:

1. They are too small. (E.g. subatomic particles)
2. They are too big/ complex for us to observe more than a small part of at any one time (E.g. economies)
3. They exist in an inaccessible region of space (This could mean far away, e.g. distant galaxies, or in hard to access places, e.g. the bottom of a deep ocean, or inside the Earth's core)
4. They existed long ago but no longer exist. (E.g. Dinosaurs)
5. They do not yet exist/ have not yet happened. (E.g. climate changes over the next 20 years)
6. They are very extended in time, meaning we can only observe a small temporal part of a longer process. (E.g. geological processes)

These difficulties in observing our subject matter all create challenges, and mean that we need to infer more than we are able to observe. For example we have to infer evolutionary lineages from current morphology, gene sequences and the glimpses we get from the fossil record, we have to infer things about the state of a whole economy from the snapshot of it we are able to observe, and so on. This reliance on inference beyond what we can observe opens up greater room for scepticism than exists for questions we can observe the answer to in quite a straightforward way.

However, the degree of access we can get to our subject matter, and the solidity of the inferences we can draw from our observations, is very variable across the different sorts of inaccessible subject matter listed above.

If we look at the examples Stanford gives of scientific claims he thinks concern "inaccessible domains of nature" (2006 p.5) we find quite a diversity of different claims. He quotes science textbook claims about ionic bonds, proliferation and homeostasis in cells, the existence of a shared ancestor for all animals, and the inflation rate of the universe shortly after the beginning of the universe.

If we try and categorise these different claims according to the list I have given above, we find several different sorts of inaccessibility. The difficulties facing us when deciding whether the animal kingdom is monophyletic is that this ancestor is no longer around, and the lineages between current species and this common ancestor are very extended in time. The issues here are not primarily to do with the objects of our investigation being too small, yet those are the issues we face with claims about ionic bonds and homeostasis in cells.

The differences between these questions are epistemically relevant – the barriers to knowledge we face are different depending on what exactly makes a subject matter inaccessible. In the case of the very small and very far away, mechanical aides to observation can be constructed to render accessible what was previously inaccessible. For example, plant cells or bacteria could not be

observed prior to the development of microscopes, and the moons of Jupiter could not be observed prior to the development of telescopes. No such mechanical aides to observation can be constructed to help with historical questions however.

That such devices are capable of rendering things observable is disputed by constructive empiricists. Van Fraassen (1980) argues that we should not commit to the literal truth of theories that go beyond what is observable, where what he means by 'observable' is that something can be observed without artificial assistance (Monton and Mohler 2021). He defines observability by saying "X is observable if there are circumstances which are such that, if X is present to us under those circumstances, then we observe it" (van Fraassen 1980). Van Fraassen's understanding of observable is thus observable 'in principle' rather than in practice, and 'in principle' observable is a very broad category. For example, Vickers notes that van Fraassen is happy to consider dinosaurs 'observable' since if a human was in the same place and time as a dinosaur they would observe them (Vickers 2022 chapter 2). Many objections to constructive empiricism have focused on the importance and coherence of the distinction between observable and unobservable (Monton and Mohler section 3.3 to 3.8). Vickers writes that this emphasis on naked eye observation is a "very fringe view nowadays, and even van Fraassen has left it behind" (ibid). More recent work by van Fraassen holds that science itself is the arbiter of what is observable (van Fraassen 2001), which seems to move his conception away from the original sense of unobservable by unaided perception.

'Inaccessible' in the sense used by Stanford (2006) seems significantly broader than early van Fraassen's concept of 'observable' since it includes things that are only unobservable in practice, i.e. things van Fraassen is happy to consider observable. However it is hard not to think that this breadth is too great, as the sorts of evidence and inference we have to rely on is so different between these different areas. A better way to frame an anti-realist argument would be to talk about a particular sort of inaccessibility, rather than accessibility in general.

Returning to the issue of historical questions, the reliability or otherwise of mechanical aides to observation is irrelevant to help us find out about the distant past. There we have to look at the traces left behind by past events and then infer what sorts of events might have left those traces. Eliminative inference clearly plays some role here, however in any actual case we will usually also have important forms of evidence that are not eliminative in character. For instance, Stanford himself discusses experimental taphonomy in which scientists recreate fossilisation in laboratories in a way that mirrors how fossilisation occurs in nature (Stanford 2011). Once we have understood how these processes work we can project this knowledge onto the past as an account of how natural fossils came to be. Stanford argues that this shows the "heterogeneity of forms of evidence,

inference and argument in scientific contexts”, a point that sits well with the positions I took in chapters 5 and 6. However, the sort of projective evidence Stanford points to in this case is not always available, as Vickers (2022) notes. Applying the issue to the different sorts of inaccessibility I discuss above, it is harder to see how projective evidence could be brought to bear on questions that are inaccessible because their subject matter is too small than to see how projective evidence could be brought to bear on questions about subject matters that are inaccessible either because they happened in the past or are too complex, big or temporally extended to be fully accessible. Different sorts of epistemic strategy are available depending on what precise sort of inaccessibility we are dealing with, and the worries we might have about these strategies varies across different sorts of difficulty.

A final point to make about inaccessibility is that some types of inaccessibility may be temporary, while others seem more permanent. Putting it more generally, the degree to which a subject matter is inaccessible is liable to change over time. For example, entire economies are not fully accessible because they are too big and complex to be fully observed. However over the course of human history we have developed more extensive bureaucracies to record economic transactions, the capacity to store and analyse larger amounts of data through computers and so on. This means that the degree of access we can get to an economy as a whole is much greater now than it was in, say, 1700 AD or even 1950 AD. Further examples of this variability in accessibility I have already mentioned in connection with microscopes and telescopes.

So the fact that many areas of science deal with inaccessible subject matters hides a great deal of diversity in what that inaccessibility consists in, and the means by which we might get around the subject matter’s inaccessibility. Whether any of these means of getting around inaccessibility succeed is largely independent of whether any of the others succeed. For example, whether an electron microscope is reliable for the purposes it is put to is entirely independent of whether randomised control trials are an effective means of discovering the efficacy of a drug. Both of these questions are also independent of whether we can piece together the history of life using genetic analysis. There is no interesting commonality here between the topics that would mean we could generalise about the unreliability of a method for gathering knowledge in one area from the fact the method is unreliable in another area.

What we see instead is many different kinds of accessibility or inaccessibility, each with a different profile of epistemic concerns and possibilities. Paying attention to the different kinds of inaccessibility and trying to work out what epistemic position each sort puts us in is important. Nevertheless, this is not to say that we cannot identify a manageable number of types of

inaccessibility in science. We do not necessarily need to be particularists when it comes to individuating contexts in a way relevant to grounds for scepticism in those contexts.

Section 2.2: Unfamiliarity

Another way of attempting to delineate the limits of the anti-realist's scepticism is mentioned by Stanford when he talks about areas of science "far removed from human experience" (Stanford forthcoming). A source of difficulty in science is that our subject matter is often unfamiliar – its nature and behaviour is very different from the sorts of 'medium sized dry goods' we are used to dealing with, e.g. tables and chairs and so on. Some sciences concern processes very different from those we see around us. Because these things are unfamiliar, we may be less able to reason inductively about them, less able to conceive all relevant explanations of their behaviour and less able to correctly interpret data concerning them. We may also be less well equipped to give scientific theories in areas like these a sensible metaphysical interpretation of the sort needed to hold realist commitment to those theories.

As examples of such areas of science Stanford mentions particle physics and cosmology, while ecology and geology Stanford sees as less unfamiliar. Vickers (2022 chapter 6) objects to this way of characterising the proper scope of scientific anti-realism on the grounds that some parts of geology do concern things 'far removed from human experience' for example "the inner/ outer core of the earth". As a further example of a topic where unfamiliarity does not seem to present an insurmountable problem Vickers also mentions evolutionary biology where "the evolution of human beings from fish over hundreds of millions of years is most definitely 'far removed from ordinary human experience'".

Why is there a difference in the degree to which these examples of unfamiliar subject matters make it more difficult to justify belief in the literal truth of scientific theories? What is the difference between the fundamental physics situation and the situation in geology and biology? Vickers discusses the idea of a 'concept application problem': "The idea is to remain sceptical whenever there is a serious question mark over the legitimacy of applying human concepts to the domain in question." For example when it comes to quantum phenomena this would "make sense of the fact that there are so many bizarre paradoxes and radically unintuitive ideas in the field of quantum theory... One reasonable explanation of our interminable perplexity is that we just don't have adequate concepts for this context".

This suggests a further working out of the idea of unfamiliarity as a barrier to knowledge that focuses on a more fundamental sort of unfamiliarity than just the fact a phenomena is not one we have seen

before. Unfamiliarity of that sort we would expect to lessen as we spent more time studying the phenomena in question, and also would allow for more 'projective evidence' where we project from things we are familiar with to times and places less familiar. In the case of quantum physics however it is the fact that our concepts are not 'fit for purpose' that creates the issue, not just that quantum phenomena are 'far removed from human experience'.

One thing to note about the idea of a concept application problem is that it seems to narrow the scope of the anti-realist's focus considerably more than a more 'bare' notion of unfamiliarity. There are not that many areas of science where we are likely to think we face a concept application problem, and all seem to come from physics.

Problems Vickers raises for the idea of a concept application problem is that it is hard for us to know when we face one, and judgements on this matter are likely to be fallible even for the entire scientific community collectively, even if we could determine what those collective judgments were. The problem this raises is that if we attempt to use the idea of a concept application to decide where we ought to be anti-realists we will not know which particular areas of science it is that we should be anti-realists about, since we will not know where a concept application problem does or does not apply. However, even if we are not perfect judges of whether a concept application problem applies, it seems plausible that we will have some idea of where there might be a concept application problem and it may be reasonable to downgrade our confidence in the reliability of particular methods in those instances if those methods led scientists astray for other cases where we seemed to face similar difficulties in the past.

Section 2.3: Conceptual difficulties

Another source of difficulty in science comes from the conceptual complexity of our theories. Because scientific theories are often complex, teasing out all of their implications can be a difficult job and can take many years of scientific work. What I mean here is different to the sort of unfamiliarity or concept application problem discussed above. The issue here is not that our concepts are unsuitable, but that they are sufficiently complicated in themselves and in their interactions with other concepts that it takes a good deal of time and work to fully realise how our theories in different areas of science relate to one another, or to realise what all the empirical implications of our theories are.

A good example of this is the modern synthesis between natural selection and Mendelian genetics. In the late nineteenth century and early twentieth century it was thought that Mendelian genetics and natural selection were incompatible as natural selection required a continuous range of

variation to operate on while Mendelian genetics seemed to suggest that the inheritance of traits is discontinuous. However, slowly the conceptual foundations for their reconciliation were laid by work such as that of Nilsson-Ehle. Larson explains that Nilsson-Ehle calculated that “if ten different genetic factors affect a trait, then sixty thousand variations of it might exist” (Larson 2006 p.171).

An even better example of conceptual complexity affecting scientific decisions can probably be found in taxonomy. Resolving debates about how to categorise living things requires conceptual work to figure out which species concept we ought to be working with. There is a large literature on this in philosophy of biology.

However, it is hard to see that this could form the basis for any general conclusion of the unknowability of any area of science. Not all areas of science are equally conceptually complex, and there is often good reason to think that what conceptual complexity does exist can be overcome, as the eventual synthesis between selectionism and Mendelian inheritance shows.

Is there any easy way to identify areas of science where conceptual difficulties should lessen our confidence in the reliability of a method? Well, one hallmark of such areas may be the presence of apparently intractable disputes between different groups of scientists, such as existed between biometricians and Mendelians in the early twentieth century. Such fundamental disagreements on the right way to approach a similar general topic may suggest that our understanding of the conceptual space requires further development. However, this is only meant as a tentative suggestion. I am not aware of any anti-realist argument that specifically focuses on this kind of difficulty, it just seems worth mentioning when attempting to catalogue different varieties of difficulty in science.

Section 2.4: Inconstant and Interactive Subject Matter

Another source of difficulty in some areas of science is that the subject matter has a habit of changing over time, and sometimes this change is brought about by the very act of studying a phenomenon.

The most striking examples of interactive subject matters are to be found in the human sciences. Hacking details how cases of Multiple Personality Disorder (MPD) rose rapidly in response to the diagnosis becoming a widely recognised condition. Hacking uses this example to argue for the existence of interactive kinds, that is, categories of things that change in response to our act of categorising them (Khalidi 2010). Any way of characterising people is likely to exhibit this kind of interactivity as people react to categorisation either by conforming more closely to the category type they identify with, or reacting against it by ceasing from behaviours identified with the type.

Economics provides another example of interactive subject matter. Many microeconomic decisions depend on people's perceptions of the macroeconomic picture at the time. For example, whether a firm decides to invest is based partly on how well the firm thinks the wider economy will do over the next year. Whether firms invest or not in turn is a causal factor on whether economic growth occurs. Our projections in economics therefore have the capacity to be self-fulfilling prophecies.

Section 2.5: Types of Question

Aside from aspects of the subject matter of science that can make its questions difficult to answer, the form of the question we are asking can present different sorts of difficulties. Consider the following types of scientific question:

- Forecast questions (how much will global temperatures rise in the next 20 years, how much will GDP increase by in the next quarter)
- Historical questions (how did modern humans evolve)
- Law of nature questions (can anything go faster than speed of light?)
- Current states of affairs (e.g. what is the R-number in the UK for COVID-19, what are the voting intentions this week of people aged 30-50 in the English midlands, what is the temperature on the surface of Venus)
- Particular causal explanations for events (e.g. what caused extinction of dinosaurs, what caused UK to be particularly badly affected in first wave of COVID pandemic)
- Causal mechanism questions for types of event (e.g. what causes bowel cancer, what causes high inflation)
- Explanations of trends (e.g. falling crime rate, rising global temperatures)

All of these questions are often found in scientific enquiry, however at first glance some seem likely to be harder to answer than others. Or at the very least, it seems likely that the difficulties involved in answering them will not be the same.

For example, current state of affairs questions are often likely to be easier than particular causal explanation of event questions. This is because answering them just involves saying how things *are* – it just needs a straightforward description. By contrast, saying what caused an event requires not just saying what happened, but also presenting evidence about causes. This seems like an extra layer of complication. Even more so, law of nature questions seem to call for a particularly high bar for evidence before we believe we have uncovered a law of nature.

Section 2.6: Analysis of Variety of Difficulties Found in Science

A theme throughout the above catalogue of types of difficulty in science has been the variability of difficulty in science. Several of the five general categories I mention above seem to sub-divide further, meaning that there are quite a large number of different sorts of difficulty to be aware of when talking about the scope of a methods reliability. In the next section I will attempt to make use of these different categories of difficulty to argue that Stanford's new induction on the history of science should be broken up into many different such inductions, and the strength of each evaluated separately. Similar points would stand for any attempt to assess a methodological tool in science against the history of science. We should generalise only within parameters set in terms of a type of difficulty.

One of the reasons why this is necessary is that each of the sorts of difficulty discussed above seem largely independent of one another. The presence of one sort of difficulty in an area allows for little generalisation about what other sorts may also be present. Further, there is little reason to think that one type of difficulty being insurmountable (at least by the method we are considering) means that other types of difficulty will also be insurmountable.

Section 3: Relativising Historical Assessments of Methods to Types of Difficulty

Section 3.1: Limiting the Scope of Sceptical Arguments

As I discussed earlier, the two characterisations Stanford gives of the areas where he thinks scepticism is warranted seem to lump together importantly different sorts of difficulty in science. Neither inaccessibility nor unfamiliarity seem to pick out very uniform categories in terms of what barriers to knowledge we face or what different sorts of methods might be suitable to overcoming these difficulties. The question is whether this variability between types of difficulty is relevant to how reliable we can expect our use of eliminative inference to be.

There are two ways in which differences between types of difficulty could be relevant to the likely reliability of eliminative inference. First, what other kinds of evidence are available to supplement eliminative inference varies depending on what kind of difficulty we are facing. We saw this when discussing Stanford's points about experimental taphonomy. However, I have already said in earlier chapters that historical assessments of methods should be confined to a single method, so this point is not necessarily relevant. What we are interested in is how far we can generalise about the reliability of eliminative inference specifically, not the extent to which the shortcomings in eliminative inference can be supplemented.

The more important issue is whether there are important differences in our ability to conceive all relevant alternatives in topics showing different sorts of difficulty. The answer here seems to be yes.

The place where we would really expect the problem of unconceived alternatives to be a particular issue is with areas where we face a concept application problem of the sort described above. Stanford himself seems to agree that this is an area where the problem of unconceived alternatives is stronger than in other places. We can perhaps be a bit more general here and say that any area of science where we face a high degree of unfamiliarity may be more susceptible to the problem of unconceived alternatives. Where an area of science deals with objects, phenomena and processes very unlike anything else we have experienced then our knowledge of the range of possible explanations is more likely to be incomplete. This is true regardless of whether the kind of unfamiliarity we are dealing with is more transient because an area has only recently begun to be investigated or more fundamental as seems to be the case in quantum physics.

However, even if it is true that very unfamiliar areas of nature may be poor places to apply eliminative reasoning, this point does not generalise over to any of the other kinds of difficulty. It may be that inaccessible subject matters also pose a particular issue of unconceived alternatives, but that seems to require separate argument, and on the whole the idea seems less plausible here. For example, we do not have access to the origins of modern species in terms of direct observation. These events happened over too long a timescale too long ago. However, we can be confident that 'modern species came into being in their current form' and 'modern species developed from different earlier forms' really are the only two options. They collectively exhaust the logical space. This remains true when we move to more specific biological questions. For example, Stanford uses as an example of a scientific claim about an otherwise inaccessible domain of nature the claim that the animal kingdom is monophyletic. Here again, we can pretty confident of exhausting the logical space. Either the animal kingdom is monophyletic or isn't. There is no possibility of an unconceived alternative here.

The reason why we can be so confident of exhausting the logical space for these questions is perhaps that historical questions concerning evolutionary lineages are not unfamiliar in the same way as quantum mechanics. It is possible to project our understanding of inheritance and so on from the present to the distant past because there seems little reason to assume that the nature of sexual reproduction has changed dramatically in the intervening time. Familiarity seems to be the issue when it comes to conceiving all alternatives, rather than accessibility.

However, in *Exceeding Our Grasp* Stanford does actually draw his case studies from the history of our understanding of inheritance. He cites Darwin in relation to his theory of pangenesis, Weismann's theory of the germ plasm and Galton's stirp theory to show cases of scientists who could not imagine any alternative to their favoured theory of inheritance. The obvious objection

here however is that what is true of a single scientist is not necessarily true of a scientific community as Ruhmkorff argues (Ruhmkorff 2011). None of the theories listed by Stanford ever achieved widespread acceptance by the scientific community.

In any case, the theories discussed by Stanford come very early in the history of theorising about the mechanisms of inheritance. At this time it is reasonable to think that unfamiliarity of the subject may have been a factor, but that this unfamiliarity was more transient than that seen in fundamental physics. Even if we grant that exhausting the logical space was difficult for this topic at this time, the relevant feature of the topic was not necessarily inaccessibility.

A further argument against thinking that inaccessibility is what is important in terms of where eliminative reasoning is likely to be effective can be gleaned from Achinstein's analysis of Perrin's case for the existence of atoms (Achinstein 2002). There Achinstein argues at length that Perrin's claims to have exhausted all possible explanations of Brownian motion apart from the atomic theory were well grounded, despite the objects of this investigation being inaccessible by virtue of being too small.

Section 3.2: Advantages of Relativising Historical Assessments of Methods to Types of Difficulty

So far we have seen that the types of difficulty found in science are diverse, and the degree of this diversity has not always been adequately appreciated. I have talked so far about how we need to be clear on what kind of difficulty is in play when deciding on the scope of sceptical arguments in science. However, more generally it makes sense to make any evaluation of the reliability of a method relative to one of the kinds of difficulty discussed above. Whatever method we are considering, we want to know what sorts of situations it applies in, and what sorts of situations it does not apply in. The most sensible way to tell that story is to talk about what kinds of difficulty we face in different situations, and whether those difficulties are such as to render the method in question unreliable. As I have already argued, the types of difficulty discussed are largely independent of one another, so whether a method is reliable relative to one kind of difficulty is independent of whether it is reliable relative to another.

Remembering that a historical assessment of methods should be about a single methodological tool, this now gives us a picture where any lesson to be learnt from the history of science should be about the reliability of one method for topics showing a particular type of difficulty. This gives the following picture of the historical debate over scientific realism. If we can show that a particular scientific theory has been formed through methods that the history of science tells us are reliable for the

kinds of difficulty we face in the area that scientific theory is found in, then realism about that area is tenable. If the history of science tells us that a method which the theory crucially relies upon for its justification is unreliable in the presence of a type of difficulty faced in the area the theory is found in, then realism about that theory is not justified. I will now try to set out some general advantages I think this approach might have over some other ways of approaching the realism debate, before going on to highlight some challenges and how they might be overcome.

As I see it my approach has four main advantages. First, it places the proper unit of evaluation in the realism debate at a broad enough level to leave some hope for learning important epistemic lessons from the history of science. Second, it sets the limits for historical assessments of methods in a principled way. Third, it does these things in a way that I think undercuts at least some of the motivation for more particularist views of scientific confirmation. Fourth, relativising judgements about methods to a type of difficulty avoids the issues of circularity and regress that can threaten empirical meta-methodology in general (these issues were discussed in chapter 4).

With regard to the first of these points, the scope of historical arguments in the scientific realism debate is undoubtedly much narrower than it has usually been conceived, at least until relatively recently. However, there are still similarities in methodological resources and types of difficulty in science broad enough for judgements of the reliability of methods relative to types of difficulty to have some kind of historical record against which to be judged. We can still hope to learn lessons with some reasonable degree of generality from the history of science.

As far as the second advantage I have claimed goes, the reason why I think my approach is generally well motivated epistemically is that it mirrors existing attempts by anti-realists to keep their scepticism within appropriate limits. The reasons for being sceptical about only some areas of science seem like they must always have something to do with the difficulty of that topic, so it seems we should really be talking about only one kind of difficulty when characterising that topic. This builds on and makes more rigorous existing justifications for being sceptical about a method of belief formation in some places and not others.

The picture I have ended up with seems to undercut at least some of the motivation for particularism about scientific confirmation because it is able to account for the need to assess each theory on its own merits but without giving up on learning epistemic lessons from the history of science. Particularist views about confirmation emphasise the diversity of science, and often draw on case studies showing that the strength of the case for a particular scientific view cannot be easily captured by any general formula that the realist might come up with. Such case studies generally give the impression that it is not possible to paint with a very broad brush when we are trying to

identify features of theories that we should or should not believe. The case for any scientific theory is usually very diverse, featuring many different kinds of evidence, and what makes the case for any theory good or bad is often dependent on the particular subject matter at hand in a way that resists characterisation in terms of any abstract schema.

I think that my approach is able to accommodate these points in a way that still leaves room for general lessons to be learnt from the history of science, lessons which can ultimately be used as tools to defend realism or anti-realism about a given theory. It is still a consequence of my view that every scientific theory will need some degree of independent evaluation, in that we will need to look at each theory individually to determine what methods are in play, and what sorts of difficulties are present.

This is because there are usually several different methodological tools and types of difficulty simultaneously at play in any real scientific context. Settling what these all are will often require quite a substantial degree of local knowledge. This means that there is no simple step from the reliability or otherwise of some general methodological resources relative to types of difficulty, to any conclusion at all about the likely truth or falsity of any scientific theory. The following form of argument for example does not work:

1. Method A is shown by the history of science to be unreliable in the presence of difficulty D
2. Method A was used to form theory T
3. D is present in the topic T is about
4. Therefore T is probably not true

The reason this does not work is it neglects the fact that there may be other methods in play when it comes to justifying T. In order to supplement this argument so it does work we need to include as a premise something like 'without the resources supplied by A the overall case for T is insufficient'. However establishing that premise to be true requires an intimate knowledge of the overall evidence for T, and so requires evaluating a theory individually.

I have intentionally left a degree of vagueness in what might count as a method or a topic. The main point I want to make is that when we are talking about historical evidence for anti-realist positions in the realism debate we need to ask 'what method are we talking about?' and 'which topics does this historical evidence show that method to be unreliable for?'. How we answer those questions strikes me as likely to be context dependent – i.e., how broadly or narrowly we draw methods or topics will vary by exactly what argument we are trying to make using these categories. This means, again, that determining whether the methods and difficulties of two different theories are appropriately similar

for the falsity of one to undermine our confidence in the other will not be easy. The extent to which the suitability of realist commitment to any given theory will require evaluation that goes beyond any broad lessons is not trivial.

If the situation is this complicated, does that not undermine the possibility of learning any meaningful epistemic lessons from the history of science? In chapter 4 I spoke about the problem posed to historical assessments of method by confirmatory holism. A similar worry seems to arise here. If we are supposed to assess a method by asking whether following it has generally led to true theories in the past, then it seems that what we should be looking for is cases where following that method led to theories we now know to be false. However, what do we make of the situation if several different types of difficulty were present in such a case? Which type of difficulty should we say that the method is unreliable in the presence of? We might also worry that even if we find no cases of a method being unreliable relative to a kind of difficulty this is not because of the reliability of that method, but for the coincidence that its shortcomings were covered over by the presence of some more reliable method in each of case. There seems to be no clear route from the historical record to conclusions about the reliability or otherwise of a method relative to a type of difficulty.

The situation could be complicated even further if either methodological resources or types of difficulty show interaction effects. What I mean here is that sometimes it could be possible that having two types of difficulty present together could create issues that are not created by either type of difficulty on its own. This would mean that the appropriate level to relativise judgements about a method to would become both kinds of difficulty when presented as a pair, rather than either individually.

I don't have a general response to this kind of worry. Instead I can only offer the hope that careful analysis of the history of science will sometimes be able to cut through such difficulties to extract convincing epistemological lessons. Such lessons will not be possible in general to derive from only a handful of case studies however. The more cases we have where a particular method has led to success or failure in the presence of a particular kind of difficulty, despite other sorts of difference between each case, the more confident we can be of having established something about the reliability of that method in different circumstances.

I suppose we might also draw an analogy between the situation we face in empirical meta-methodology and the situation we face in many empirical sciences. Isolating different causal factors in complex systems in order to come up with workable general principles is generally difficult to do, but that does not mean that there are no worthwhile lessons of any generality to be gained by the attempt to understand these systems. Similarly, finding that it is difficult to extract general epistemic

lessons from the history of science should not necessarily dissuade us from trying. If we find that the causal principles of a field are complicated that does not lead us to conclude that the causes of each separate event are so particular that generalisation is impossible. It just teaches us to be more nuanced in our search for general principles. This should be our attitude with scientific method – not that methods are so context sensitive that little can be said about their general reliability or unreliability, but that methodology is sufficiently complicated that we should expect some nuance in these judgments and should not be overly ambitious when deciding how wide their scope might be.

Returning finally to the fourth advantage I claimed for my approach to the realism debate, relativising judgments of methods to a type of difficulty means that issues around circularity and regress in empirical meta-methodology can be avoided. I spoke in chapter 4 about how if we use a method in our empirical assessment of that method then there seems to be a circularity in this that undermines any claim we might make to have given the method further support through our empirical assessment. Once we realise that judgements about the reliability of a method are topic relative, this worry can perhaps be avoided so long as the types of difficulty we face in the meta-methodological context are different from the ones we are trying to assess the reliability of that method for. We might still worry about the possibility of a regress here – how are we to justify the method we are using relative to the meta-methodological context? However this seems like a danger however we justify anything – it is not clear any more that this is a problem specific to empirical meta-methodology.

Conclusion

In this chapter I have argued that when assessing methods we should be clear on which topics we think those methods are and are not reliable for. The best way to individuate topics is in terms of what kind of difficulties we face in different areas. Taken together with my arguments about what kind of methods might be amenable to historical assessment, this now gives us a picture of the right level for historical debates over scientific realism to take place. This approach has the advantage of carving out some role for learning from the history of science at the same time as recognising the need for detailed individual assessment of any particular theory we might be considering realism or anti-realism about. My hope is that this reconciles the historical scientific realism debate with at least some of the insights of methodological pluralism in general, and particularists about scientific confirmation especially. At the same time, establishing anti-realist (or realist) conclusions that are anything resembling global across science by means of a single induction on the history of science no longer looks like a realistic prospect. However, such inductions can possibly provide important tools for deciding whether to be realist or anti-realist about a given theory.

Chapter 8: Review, Comparison with Other Work in the Field, and Conclusion

Introduction

In this chapter I will review the position I have taken on historical arguments in the scientific realism debate, and the arguments I have made in support of that position. I will compare my position with some of the existing work that is relevant to localism in the scientific realism debate, to bring out the points of agreement and disagreement between myself and other writers. The writers whose positions I will compare with own are Asay, Saatsi, Vickers, Fitzpatrick and Henderson. My intention in doing this is to show further how my work fits into the wider debate. I will then go on to give examples of how historical arguments might go if they are made along the lines I suggest. I will finish by pointing to where I think future work could be done to build on what I have done in this thesis.

Section 1: Summary of Thesis Argument

The most important points of my thesis can be summarised as follows:

1. A scientific realism debate focusing on novel predictive success is poorly motivated and leaves out a good deal of important science, including but by no means limited to evolutionary biology. (See chapters 2 and 3.)
2. We should view the historical realism debate as a meta-methodological debate that attempts to justify or refute methodological principles empirically. (See chapter 4.)
3. When attempting to assess the reliability of any method of forming beliefs we should be clear about *which* method we are talking about, and try to make sure that our target really is unitary in an appropriate way. (See chapter 4.)
4. Once we realise that there is no universal scientific method it becomes clear that historical arguments cannot establish global epistemic attitudes towards science, or even a global formula for deciding our epistemic attitudes towards science. Any version of the general 'infer truth from success' idea will either be relevant to only a part of science or will conflate different methods of belief formation in an illegitimate way. (See chapter 5.)
5. Isolating methods in a well-motivated way that still leaves room for historical assessment is very challenging. However it seems we can identify some methodological resources, such as inference types for example, that are widely shared across scientific methodologies. In some cases this allows historical arguments to be made that do seem to be appropriately localised to a method. Stanford's problem of unconceived alternatives gives an example of this. (See chapters 5 and 6.)

6. It is also important to recognise that a method of forming beliefs may be reliable for some topics but not others. This means we also need some way of individuating topics in a well-motivated way. (See chapters 4, 6 and 7.)
7. Attempting to avoid historical arguments entirely by pointing to the context-sensitivity of methods does not work as the anti-realist can argue that although whether a method *really is* reliable may vary by context, we have shown ourselves to be poor judges of whether a method applies in a given context or not. (See chapter 6.)
8. To get around this issue we need a principled way to distinguish between *types* of context. The best way to do this is by categorising contexts according to the types of difficulty we face in them. Type of difficulty present is the most plausible reason why a method would be reliable in some contexts but not others. There are quite a number of distinct varieties of difficulty in science. Recognising which types of difficulty are present in a scientific context may not always be possible, but usually is (See chapter 7.)
9. It is probably not possible to give a complete list of difficulties and methods in science. There is probably a contextual element to whether a given characterisation of a method or a difficulty is well framed or not, where this in part relies on what sort of argument we are trying to make about the methods or difficulties under discussion. (See chapter 7.)
10. Because there are numerous methods and difficulties at play in any real scientific context, and identifying these requires substantial knowledge of the particular theory we are considering, there is no simple route from a historical assessment of a method relative to a type of difficulty to realism or anti-realism about any specific scientific theory. In this way individualists in the scientific realism debate are right. However this does not threaten the value of historical inductions on science as a tool for assessing the appropriateness of realism about current theories. (See chapter 7.)

I will now contrast this position with those taken by other authors whose work seems particularly relevant to mine. My aim here is not so much to argue that these other writers are wrong and I am right where we disagree. In some cases I have already argued that earlier, but even where I have not, my aim is more to illustrate further what my position is and how it fits into the existing debate. Highlighting areas of agreement and disagreement seems a good way to achieve this.

Section 2: Similarities and Differences between my Position and Other Authors

Section 2.1: Saatsi on Replacing Recipe Realism

The first philosopher whose position on the scientific realism debate I want to compare with my own is Juha Saatsi. Saatsi (2017) argues that scientific realism should be viewed as a “global meta-level

attitude” towards science which amounts to “an adherence to the No Miracles intuition according to which the impressive empirical success of science is by and large down to theories latching onto reality in ways that make them empirically successful” (Saatsi 2017 p.8). This general attitude does not translate into any simple formula or ‘recipe’ for deciding where and how far we should be realists: “The idea that this is a reasonable attitude towards science says nothing in and of itself about where exactly we should place our trust in science” (ibid) but instead “A realist gains her epistemic commitments when she applies her realist attitude locally, in a piecemeal way” (2017 p.9). A realist’s account of the success of any particular bit of science will rely on exemplars, allowing the realist to make claims like: “in a domain of science like *this*, with theories or models like *that*, empirical success in *this* sense, is (probably) accountable in *those* terms” (ibid).

Saatsi is led to give this account of scientific realism by similar motivations against ‘recipe realism’ as I expressed in earlier chapters. By ‘recipe realism’ Saatsi means the idea that it is possible to give some general formula for deciding which scientific claims realism is appropriate about, such as ‘if a theory shows sufficient novel predictive success then it is probably approximately true in its working posits’. Saatsi argues that the diversity of science means that such recipes are likely to either face areas of science where they are inapplicable or unworkable, or they are likely to be ‘contentless’, i.e., too vague to be applied.

The points of agreement between Saatsi’s position and my own are considerable. We both reject universal success-based formulas for deciding our realist commitments, and for broadly similar reasons. We both argue that a more localised version of the realism debate is needed, however we both also want to leave room for learning some epistemic lessons from the history of science. We both also emphasise the need to use relevant points of similarity across scientific contexts as the means to localise discussion. Saatsi does this by using exemplars to illustrate relevant similarity, while I do it by demanding that historical inductions in the realism debate be about the same method and the same type of difficulty. In both of our accounts however there is room for judgments of similarity and difference to be to some degree contextual and hard to fully set out in advance. Whether a particular case is suitably similar to an exemplar is hard to set out in a formula. Similarly, I claim that whether a method and a type of difficulty are well characterised is partly dependent on what argument they are being used to constrain, meaning that we cannot give a complete list of all methods and difficulties in isolation.

Points of important outright disagreement between my position and the one Saatsi takes in “Replacing recipe realism” are hard to find, however there are some slight differences. First, Saatsi’s main concern seems to be with finding a way to defend the general intuition that scientific theories

are successful because they get something right about the world. He seems happy to “bite the bullet in response to the challenge that with respect to most theories we do not really know what we are realists about” (2017 p.11). However, my hope is that the work I have done in this thesis gives us tools that we can use to decide what we are realists about – or at least, tools that may sometimes be relevant to this decision. When we are looking at a particular theory we should ask what kinds of evidence, inference and methodological resource are available, what kinds of difficulty we face, and whether there are any difficulties for which none of our methodological resources have proven reliable in the history of science. If so, there is a good case for anti-realism. However, if there are very few examples in the history of science of theories that used similar methodological tools in the presence of similar difficulties and came to a false theory, then the history of science does not provide grounds for scepticism in that case. In fact, those last two sentences actually seem to constitute a general recipe for deciding our realist commitments, so perhaps what I am advocating is a form of recipe realism, even if the recipe in question leaves a lot of gaps to be filled in for any particular context.

A final point of contact between Saatsi’s position and my own could perhaps be that I give an account of which dimensions of similarity and difference we should be looking for when we use exemplars, i.e., I argue that what is relevant in terms of similarities between an exemplar and some other case is whether the same types of difficulty are present in both, and whether both are using similar methods. This may help to constrain the use of exemplars and clarify whether we should take some particular exemplar to be relevant to whatever case we are discussing.

Section 2.2: Fine and Maddy

There is a tradition of philosophers of science who argue that the only evidence that is relevant to determining our epistemic commitments to a scientific theory is the first order evidence. Such philosophers usually make this claim in order to explicitly deny that historical arguments are relevant in the evaluation of a scientific claim – i.e., in order to deny that arguments like the pessimistic induction should threaten our belief in scientific claims. Fine (1986) and Maddy (2001) are probably the two best known examples of philosophers who hold this position.

This kind of position appears to be incompatible with the project I have been undertaking in this thesis, as I have been trying to find a way to make it plausible that arguments inspired by the history of science can be relevant to determining our attitudes to current theories.

However, I have ended up with a position that agrees with this strain of thinking insofar as I think that a detailed engagement with the evidence for an individual theory is very often needed in order

to decide whether to believe that theory – historical arguments should not have as straightforward conclusions that we should believe or disbelieve any scientific theory. Instead they should be about the reliability of methods for particular sorts of difficulties. This means that when we are attempting to evaluate the first order evidence for a scientific theory, historical inductions may be relevant to helping us decide how seriously to take some particular piece of evidence, if similar sorts of evidence have been misleading or unreliable in the presence of similar sorts of difficulty.

Vickers (2022 chapter 2 section 4) makes the point that when it comes to assessing the strength of the first order evidence for a scientific claim we are best off simply asking scientists how the weight of evidence stands, as they know a lot more about it than philosophers of science do about any particular case. This seems to leave little room for historical arguments for or against realism.

However, thinking that historical evidence can be relevant to the strength of a case for a scientific claim is not necessarily the same thing as thinking that philosophers of science are ever in a position to judge the overall strength of the evidence. What we can use historical inductions for, on my account, is saying that some particular means of justifying a scientific claim is or is not reliable for a type of difficulty. At best, this would allow us to re-evaluate the strength of this or that particular piece of evidence for a claim. It does not substitute for the full scientific work of assessing the overall weight of evidence. So in that sense, the role I am advocating for historical inductions is much more modest than that envisaged by traditional realists or anti-realists. I do not generally imagine historical arguments completely undercutting scientists' own assessments of the weight of evidence – I just imagine historians and philosophers of science sometimes being in a position to challenge some particular arguments that might be offered in support of a scientific claim.

My hope is also that these more nuanced attempts to use the history of science for sceptical purposes might also be able to present historical challenges in a way that seems less fanciful and 'arm-chair-ish' than something like the traditional pessimistic induction. Perhaps they could even be an epistemic resource that scientists themselves could actually make use of in assessing certain sorts of evidence.

Anyway, there are two fallback options for how I can present the relevance of my project for those who think that assessment of scientific theories is best left to the scientists. First, attempting to learn lessons from the history of science may still be useful as a project in epistemology – it might offer insights into the circumstances in which different sorts of evidence should be taken as confirmatory, or offering general lessons about the contextuality of evidential relations for example. Something closely related to the historically focused realism debate may still be worthwhile even if we don't expect it to offer any substantial advice on which theories we should believe. Perhaps all we

are doing here is codifying judgments that scientists are already making intuitively, rather than telling them better ways to reason about evidence. But this task can still help us understand the nature of evidence and scientific reasoning better. Secondly, my thesis could be taken as arguing for the conditional claim that *if* it is possible for historical arguments to have a bearing on our epistemic commitments in science, those historical arguments must be relativised to a method and a type of difficulty in the way I have suggested.

Section 2.3: Asay

Asay is slightly unusual among advocates of localism in that he allows for the possibility of determining our epistemic commitments collectively to some degree. That is, he allows the possibility that we might be realists or anti-realists about entire fields, rather than about individual theories only. However, as we have seen Park (2019) complain, Asay does little to lay out how we should determine the proper level at which to engage in realism debates other than to say we should make these debates as narrow or as broad as seems useful to us.

I agree with what Asay has to say about the proper level at which to engage in realism debates – I agree both that it should vary depending on what sort of argument we are making, and that sometimes it might be reasonable to evaluate fairly large parts of science (e.g. entire disciplines) collectively. Why this should be is something that I can give an account of. The important thing with historical inductions on the history of science is that they are about relevantly similar methods and types of difficulty. Sometimes it may be that there is enough common ground in terms of methods and types of difficulty throughout an entire discipline that any historical induction that affects one part of the discipline will also affect others. In this way, it might be reasonable in some circumstances to be sceptical of an entire discipline based on a historical induction relativised to the methods and difficulties of that discipline. In other cases, the methods used and the difficulties faced may be more esoteric such that there are not relevant similarities between one theory and others. In that case each theory will have to be evaluated independently.

Section 2.4: Henderson

Recall that Henderson (2018) worries that if we claim that each theory should be evaluated locally using only first order evidence we lose the ability to justify any contested piece of scientific reasoning by comparison with less contentious pieces of reasoning. Such comparisons always break down due the differences in context between the two cases being compared on the view Henderson ascribes to the localist, therefore no justification for the contentious case can be provided in this

way. The realist therefore loses perhaps their most important resource by resorting to an 'each theory on its own merits' approach.

On my view it is possible to compare a contested piece of reasoning with a more established one in order to justify the contentious one, so long as what we are comparing are genuinely similar methods in each case, and similar difficulties are also present in each case. This regains comparative resources for justifications of realism, while also constraining the scope of anti-realist arguments. My approach therefore seems to reconcile a version of localism with Henderson's worries.

Section 2.5: Summary of how my work fits into the debate

To summarise the most important point for how my work fits into the wider debate, I hope I have shown how many of the insights of localists can be reconciled with the idea that epistemic lessons can be learnt from the history of science. Even if detailed engagement with local scientific details is needed to evaluate a scientific theory, historical arguments can still be relevant to the attitudes we take to some of this evidence.

Section 3: Suggestions for Further Work

One area where I think there is certainly more to say is with regard to individuating methods for historical assessment. In chapter 5 I wrote at length of the difficulties facing attempts to individuate methods, and although I gave some indication of how this might be done in chapter 6, this is far from a complete or systematic account. One important area for further work is therefore trying to say more about how methods might be individuated, and which sorts of methods might be amenable to historical assessment of a sort that will be illuminating.

A second area for further work is in applying the framework I have been trying to develop. The main argument of this thesis has been that how we understand the roles of historical inductions on the history of science needs to be reframed. The most natural suggestion for further work leading on from this is to look again at the cases that have been raised as fodder for the pessimistic induction, and re-evaluate what lessons can be learnt from these cases. None of these should be interpreted as showing that truth cannot be inferred from success in science. Instead we should take them as evidence that the methods and inferences in play in these examples are unreliable for at least one of the types of difficulty found in them. Once we have localised which methods and difficulties such cases most plausibly tell us something about, we can possibly then use those conclusions to challenge particular aspects of current science. What judgments we are led to in these cases may also provide further means of evaluating the framework I have put forwards. If it turns out to be impossible to apply my framework in the way suggested in a plausible way, that would be a serious mark

against my framework. If however applying my framework to classic pessimistic induction cases seems to yield more plausibly constrained and illuminating lessons that cannot be easily dismissed by other sorts of localist argument, that will seem like a strong consideration in favour of the framework I have suggested.

References

- **Achinstein, Peter (2002)**. Is there a valid experimental argument for scientific realism? *Journal of Philosophy* 99 (9):470-495.
- **Alston, W. P. (1991)**. *Perceiving God: The epistemology of religious experience*. Ithaca, NY: Cornell University Press
- **Asay, Jamin (2019)**. Going local: a defense of methodological localism about scientific realism. *Synthese* 196 (2):587-609.
- **Barnes, Eric Christian (2022)**, "Prediction versus Accommodation", *The Stanford Encyclopedia of Philosophy* (Winter 2022 Edition), Edward N. Zalta & Uri Nodelman (eds.), URL = <<https://plato.stanford.edu/archives/win2022/entries/prediction-accommodation/>>.
- **Bowler, Peter J. (1985)**. Evolution: The History of an Idea. *Journal of the History of Biology* 18 (1):155-157.
- **Boyd, Richard (1983)**. On the current status of the issue of scientific realism. *Erkenntnis* 19 (1-3):45 - 90.
- **Briggs, D. and Walters, S.M. (1997)**. *Plant Variation and Evolution*. Third edition. Cambridge: Cambridge University Press.
- **Bueno, Otávio (2012)**. Styles of reasoning: A pluralist view. *Studies in History and Philosophy of Science Part A* 43 (4):657-665.
- **Chakravartty, Anjan (2017)** "Scientific Realism", *The Stanford Encyclopedia of Philosophy* (Summer 2017 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/sum2017/entries/scientific-realism/>>.
- **Chalmers, Alan (1986)**. The Galileo that Feyerabend Missed: An Improved Case Against Method. In Schuster and Yeo (eds.) *The Politics and Rhetoric of Scientific Method*. Dordrecht: D. Reidel Publishing Company.
- **Crombie, A. C. (1981)**. Philosophical perspectives and shifting interpretations of Galileo. In J. Hintikka, D. Gruender, & E. Agazzi (Eds.), *Theory change, ancient axiomatics and Galileo's methodology* (pp. 271–286). Dordrecht: Reidel.
- **Crombie, A. C. (1994)**. *Styles of scientific thinking in the European tradition: The history of argument and explanation especially in the mathematical and biomedical sciences and arts*. Duckworth.
- **Darwin, C. (2003 [1859])**. *On the Origin of Species*, 1859 (1st ed.). Routledge. <https://doi.org/10.4324/9780203509104>

- **Descartes, Rene (1998 [1641]).** *Meditations and Other Metaphysical Writings*, translated Desmond M. Clarke. Penguin Books.
- **Dupré, John (1993).** *The Disorder of Things: Metaphysical Foundations of the Disunity of Science*. Harvard University Press.
- **Feyerabend, Paul (1975).** *Against Method*. London: New Left Books.
- **Feyerabend, Paul (1978).** *Science in a Free Society*. London: New Left Books
- **Fine, Arthur (1986).** Unnatural attitudes: Realist and instrumentalist attachments to science. *Mind* 95 (378):149-179.
- **Fitzpatrick, Simon (2013).** Doing away with the No Miracles Argument. In Dennis Dieks & Vassilios Karakostas (eds.), *Recent Progress in Philosophy of Science: Perspectives and Foundational Problems*. Springer.
- **Galilei, Galileo (1630).** *Dialogue Concerning the Two Chief World Systems, Ptolemaic and Copernican*.
- **Galilei, Galileo (1638).** *Dialogues Concerning Two new sciences*.
- **Giere, Ronald N. (1984).** *Understanding Scientific Reasoning*, second edition, New York: Holt, Rinehart, and Winston. First edition 1979.
- **Glymour, Clark N. (1980).** *Theory and Evidence*, Princeton, NJ: Princeton University Press.
- **Hacking, Ian (1982).** Language, truth and reason. In Martin Hollis & Steven Lukes (eds.), *Rationality and Relativism*. MIT Press. pp. 48--66.
- **Hacking, Ian (1992).** The self-vindication of the laboratory sciences. In Andrew Pickering (ed.), *Science as Practice and Culture*. University of Chicago Press. pp. 29--64.
- **Hacking (2002).** 'Style' for historians and philosophers. In idem (2002), *Historical ontology* (pp. 178--199). Cambridge, MA: Harvard University Press. (First published 1992)
- **Hacking, Ian (2012).** 'Language, Truth and Reason' 30years later. *Studies in History and Philosophy of Science Part A* 43 (4):599-609.
- **Henderson, Leah (2018).** Global versus local arguments for realism. In Juha Saatsi (ed.), *The Routledge Handbook of Scientific Realism*. Routledge. pp. 151-163
- **Howson, Colin (1990).** "Fitting Your Theory to the Facts: Probably Not Such a Bad Thing After All", in *Scientific Theories, (Minnesota Studies in the Philosophy of Science, Vol. XIV)*, C. Wade Savage (ed.), Minneapolis: University of Minnesota Press, pp. 224--244.
- **Jenkin, Fleeming (1867).** The Origin of Species. *North British Review* 46: 277-318

- **Johannsen, Wilhelm (1955 [1909]).** Concerning Heredity in Populations and in Pure Lines. Translated by Harold Gall and Elga Putsch. In *Selected Readings in Biology for Natural Sciences*. Chicago: University of Chicago Press. Pp. 172-215.
- **Khalidi, Muhammad Ali (2010).** Interactive kinds. *British Journal for the Philosophy of Science* 61 (2):335-360.
- **Kitcher, Philip (1993).** *The Advancement of Science: Science Without Legend, Objectivity without Illusions*, Oxford University Press.
- **Kuhn, Thomas S. (1962).** *The Structure of Scientific Revolutions*. University of Chicago Press.
- **Kusch, Martin (2010).** Hacking's historical epistemology: a critique of styles of reasoning. *Studies in History and Philosophy of Science Part A* 41 (2):158-173.
- **Larson, Edward J. (2004).** *Evolution: The Remarkable History of a Scientific Theory*. Modern Library.
- **Laudan, Larry (1981).** A confutation of convergent realism. *Philosophy of Science* 48 (1):19-49.
- **Laudan, Larry (1984).** *Science and Values: The Aims of Science and Their Role in Scientific Debate*. Berkeley, CA: University of California Press.
- **Laudan, Larry (1996).** *Beyond Positivism and Relativism: Theory, Method and Evidence*. Boulder, CO: Westview
- **Maddy, Penelope (2001).** Naturalism: Friends and Foes. *Noûs* 35 (s15):37-67.
- **Magnus, P. D. & Callender, Craig (2004).** Realist Ennui and the Base Rate Fallacy. *Philosophy of Science* 71 (3):320-338.
- **Mill, John Stuart (1916 [1872]).** *A System of Logic: Ratiocinative and Inductive: Being a Connected View of the Principles of Evidence and the Methods of Scientific Investigation*, 8th ed. London: Longman, Green, and Co.
- **Monton, Bradley and Mohler, Chad (2021).** "Constructive Empiricism", *The Stanford Encyclopedia of Philosophy* (Summer 2021 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/sum2021/entries/constructive-empiricism/>>.
- **Musgrave, A. (1985).** Realism versus constructive empiricism. In P. M. Churchland & C. A. Hooker (Eds.), *Images of Science* (pp. 197–221). Chicago, IL: Chicago University Press
- **Nola, Robert & Sankey, Howard (2007).** *Theories of Scientific Method: An Introduction*. Acumen Publishing.
- **Norton, John D. (2003).** A material theory of induction. *Philosophy of Science* 70 (4):647-670.
- **Park, Seungbae (2019).** Localism vs. Individualism for the Scientific Realism Debate. *Philosophical Papers* 48 (3):359-377.

- **Popper, Karl (1959).** *The Logic of Scientific Discovery*, London: Routledge, 2002
- **Popper, Karl (1963).** *Conjectures and Refutations*, London: Routledge, 2002.
- **Psillos, Stathis (1999).** *Scientific Realism: How Science Tracks Truth*, Routledge.
- **Psillos S (2007).** The fine structure of Inference to the Best Explanation. *Philosophy and Phenomenological Research* 74: 441-448.
- **Putnam, Hilary (1975).** *Mathematics, Matter and Method*. Cambridge University Press.
- **Quine, Willard (1992).** *Pursuit of Truth*, rev. edn. Cambridge, MA: Harvard University Press
- **Rescher, Nicholas (1977).** *Methodological Pragmatism*. Oxford: Blackwell.
- **Ruhmkorff, Samuel (2011).** Some Difficulties for the Problem of Unconceived Alternatives. *Philosophy of Science* 78 (5):875-886.
- **Russell, Bertrand (1912).** *The Problems of Philosophy*. Barnes & Noble.
- **Saatsi, Juha (2009).** Form vs. Content-driven Arguments for Realism. In P. D. Magnus & Jacob Busch (eds.), *New Waves in Philosophy of Science*. Palgrave-Macmillan.
- **Saatsi, Juha (2017).** Replacing recipe realism. *Synthese* 194 (9):3233-3244.
- **Saatsi, J., & Vickers, P. (2011).** Miraculous Success? Inconsistency and untruth in Kirchhoff's diffraction theory. *British Journal for the Philosophy of Science*, 62(1), 29–46
- **Sankey, Howard (2000).** Methodological pluralism, normative naturalism and the realist aim of science. In Howard Sankey & Robert Nola (eds.), *After Popper, Kuhn and Feyerabend: Recent Issues in Theories of Scientific Method*. Dordrecht/Boston/London: Kluwer Academic Publishers. pp. 211-229.
- **Sankey, Howard & Nola, Robert (2000).** A Selective Survey of Theories of Scientific Method. In Robert Nola & Howard Sankey (eds.), *After Popper, Kuhn and Feyerabend: Recent Issues in Theories of Scientific Method*. Dordrecht/Boston/London: Kluwer Academic Publishers. pp. 1-65.
- **Simpson, George Gaylord (1944).** *Tempo and Mode in Evolution*. New York: Columbia University Press.
- **Stanford, P. Kyle (2003).** Pyrrhic victories for scientific realism. *Journal of Philosophy* 100 (11):553 - 572.
- **Stanford, P. Kyle (2006).** *Exceeding Our Grasp: Science, History, and the Problem of Unconceived Alternatives*. New York, US: Oxford University Press.
- **Stanford, P. Kyle (2011).** Damn the Consequences: Projective Evidence and the Heterogeneity of Scientific Confirmation. *Philosophy of Science* 78 (5):887-899.

- **Stanford, P. Kyle (2015).** Catastrophism, Uniformitarianism, and a Scientific Realism Debate That Makes a Difference. *Philosophy of Science* 82 (5):867-878.
- **Stanford, P. Kyle (2021).** Realism, Instrumentalism, Particularism: A Middle Path Forward in the Scientific Realism Debate. In Lyons and Vickers (eds.) *Contemporary Scientific Realism: The Challenge from the History of Science*. Oxford University Press.
- **Van Fraassen B (1980).** *The Scientific Image*. Oxford: Clarendon Press.
- **Vickers, Peter (2012).** Historical magic in old quantum theory? *European Journal for Philosophy of Science* 2 (1):1-19.
- **Vickers, Peter (2013).** A Confrontation of Convergent Realism. *Philosophy of Science* 80 (2):189-211.
- **Vickers, Peter (2019).** Towards a realistic success-to-truth inference for scientific realism. *Synthese* 196 (2):571-585.
- **Vickers, Peter (2022),** *Identifying Future-Proof Science* (online edition, Oxford Academic)
- **Woodward, James and Lauren Ross (2021).** "Scientific Explanation", *The Stanford Encyclopedia of Philosophy* (Summer 2021 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/sum2021/entries/scientific-explanation/>>.
- **Zahar, Elie (1983).** *Einstein's Revolution: A Study In Heuristic*, La Salle, IL: Open Court.