Design Knowledge and the Instruments of Detection

David Trembath (B.Eng. *Curtin*, BA *UWA*)

This thesis is presented for a doctorate of philosophy in the discipline of Philosophy, School of Humanities at the University of Western Australia.

November 2010

Abstract

The thesis defends a form of scientific realism based on the use of instruments of detection and operations of measurement in scientific practice (and operations more generally construed). I argue that we can have knowledge of the detectable and measurable part of the world in virtue of what our instruments of detection and measurement indicate. Although related to Ian Hacking's entity realism, the kind of realism I defend is substantially different. Unlike Hacking, I wish to argue for an important relationship between a theory and the entities over which that theory ranges. Thus, an important part of the project will be to justify those 'secure' parts of a theory used to interpret what is indicated by our instruments of detection and measurement. The thesis also considers what ontological commitments are required for this form of realism, and how these commitments relate to important developments in the realism debate within the philosophy of science.

Contents

Acknowledgements

I would like to acknowledge the considerable guidance and help provided by my principal supervisor Barry Maund. It was originally the philosophy of science course that Barry coordinated that sparked my interest in this area. As an engineer I did not know that such questioning of the scientific orthodoxy was allowed, let alone justifiable. Indeed, it is something of a repair to this ‘pre-philosophy of science’ knowledge that this project attends to.

I also would like to thank Stewart Candlish for reviewing various drafts of my thesis, and maybe more importantly, teaching me the value of some attention to detail. While it is a skill I am sure I have yet to master, taking seriously Stewart’s questioning, of what seem like small errors or innocuous equivocations, has led to important developments in my thesis. Needless to say, all remaining errors are my own.

Dr Nic Damnjanovic has also provided important help, in clarifying and simplifying the thesis, but also by offering me the chance to be the principal lecturer in the Philosophy of Science course at The University of Western Australia. Teaching this course provided me greater insight into the problems I was attending to in my thesis, while also exposing some gaps in my understanding that needed addressing.

Finally, I would like to acknowledge the help that my family and friends provided, especially Jennifer, who has been my touchstone of commonsense.

To you all, I whole heartedly thank you.

Introduction

This thesis is an attempt to justify a realist interpretation of scientific language and practice. Scientists often introduce theories and posit unobservable entities with a degree of agnosticism. Yet there are also theories and theoretical entities that the scientific community considers beyond reproach. The causes of many diseases, the atomic nature of elements, several thermodynamic laws, many of the properties of electromagnetism, are considered more or less established by both applied and pure scientists. This thesis is an attempt at providing some explanation, and more importantly justification, for this shift from agnosticism towards realism. What this realism amounts to, however, is contestable, maybe even controversial. So an important part of my thesis will be to justify my kind of realism over other important interpretations of scientific practice.

In order to test and thus strengthen my realist interpretation of scientific practice, I outline what I consider the most defensible form of anti-realism – van Fraassen's constructive empiricism. Yet the empiricism that informs this anti-realist position also features heavily in my thesis. In fact, an important part of my thesis is an attempt to extend the epistemic status that empiricists bestow on the observable to what is detectable. Not that the use of the concept of detectability is a particularly new tactic: the idea is important in the development of scientific entity realism, powerfully argued for by Ian Hacking. Hacking, however, does not say much about these entities. This silence might be justified. There is, according to some, a long history of scientific theories that, having once been successful, in the end turned out to be false. If this is right, certain difficult issues can be raised regarding the realist's view of science. For instance, some realists explain the success of science by appealing to the approximate truth of the theories involved. However, if there is a long history of successful theories that turn out to be false, and false in a way that they are not even approximately true, these realists have lost what they hoped to use as an explanation. Entity realism can avoid these issues, but at the price of not saying much about what our current best science claims exists.

In chapter 1 I set out some of the influences that inform the type of realism I wish to defend. Apart from Hacking and van Fraassen, my thesis draws upon some of the insights offered by the early American pragmatists, Charles Peirce and John Dewey. Peirce attempted to *fix* or *secure* beliefs regarding a scientific theory in virtue of both the scientific method and some 'external permanency'. I too wish to provide some security to important parts of certain scientific theories. However, Peirce's pragmatism led to a complex and controversial ontology, the likes of which I hope to avoid. Dewey attempted to bypass the metaphysical difficulties Peirce found himself in by focussing on the applied sciences and the methods they use to solve problems. Similarly, throughout much of the thesis I will draw upon the practices and theories of the applied sciences in order to find a way forward.

I begin analysing the applied sciences in earnest in chapter 2. Here I argue that, as a result of what I term an ‘informative test’, we can have secure theory. This idea conflicts with what it is to ‘genuinely test’ a theory as developed by Karl Popper. I argue that, in so far as Popper’s falsificationism is meant to demarcate science from non-science, his account depends upon something like an informative test. Without this second type of test, Popper cannot demarcate applied from pure science. If the applied sciences are to be considered scientific, and distinct from the pure sciences, then they need a separate demarcation criterion. Yet Popper thought that the only form of genuine testing was a severe test. If this were right, the applied sciences are either non-scientific or a form of pure science. Either option, I argue, is unacceptable.

Having argued that we require secure theory for a complete explanation of scientific success, in chapter 3 I provide an account of how we might achieve this. Central to the account is what has been termed the ‘differential approach’ to operationalism. Although operationalism has been the subject of serious criticism, the differential approach provides considerable insight as to how we might test two related and operationally ‘defined’ theories against one another. If the theories survive such a test, this justifies treating the theories as secure. This is important information to have when considering those theories used to interpret the operation of our instruments of measurement and detection. All this, however, will seem rather tangential to the goals of constructive empiricism. A constructive empiricist argues that observability is the important epistemic criterion, especially when judging the empirical adequacy of a theory. However, I argue that observation, properly construed, has analogous features to detection. Indeed, something like the differential approach is used to justify claims regarding what is observed of the observable world. If this is right we can extend what is empirical to include the detectable world.

Having extended, by analogy, what counts as empirical, chapter 4 considers what we can count as real on the basis of empirical information. Here I introduce what I term a *structure of properties*. Roughly, a structure of properties is that combination of properties that constitute that part of an entity (or entities) that we can detect and measure, the results of which warrant the belief that an important part of a theory is true (or approximately so). An epistemic commitment to such structures I term ‘differential realism’. Much of the chapter is devoted to clarifying the relationship between a structure of properties, our instruments of detection and measurement and the relevant parts of a theory. Having a clear understanding of what differential realism amounts to will be important when we come to consider other similar forms of realism.

Having established something of the epistemology and ontology of the form of realism I am developing, I next consider, in chapter 5, accounts similar to mine. Like differential realism, these forms of realism consider what might be preserved as science progresses. An important factor in the development of most forms of preservative realism is what is termed the pessimistic meta-induction (PMI). As mentioned above, many successful theories have turned out to be false. This is not just to note the trivial point that, in so far as a theory might be approximately true it is, strictly speaking, false. Rather, the PMI is meant to show that most successful theories failed to genuinely refer to what they took to exist, and thus can't even be approximately true (at least according to some). If the PMI is right, we have good inductive grounds for believing that those theories used to describe the detectable world will turn out to be false. It is not surprising that the realists who take the PMI seriously provide rather weak forms of realism. And yet, I argue, the basis for the PMI has not been established. This undermines one of the central motives for being conservative in what we take as real. I then argue that differential realism avoids criticism that is levelled at its close relatives – deployment realism and semi-realism.

The final chapter considers the development and use of the Ramsey sentence in understanding, or accounting for our realist intuitions about scientific theories. After providing a brief outline of its development, I consider the recent use of the Ramsey sentence by the structural realist. Various criticisms have been levelled at forms of structural realism based on this approach. I leave these usual problems to one side so as to consider a tension in the goals of structural realism and their use of the Ramsey sentence approach. I argue that the use of additional machinery by the structural realist to avoid this tension makes it difficult to assess what additional insight the Ramsification of a theory is supposed to add. I also consider the use of the Ramsey sentence (or an important version of it) in sorting out issues of reference and approximate truth in the scientific realism debate. I argue that the success of the Ramsey sentence approach in this regard also depends on factors or criteria independent of such an approach. For instance, differential realism can provide a guide as to what counts as the empirical success of a theory. This is important as it is the approximate truth of a Ramsey sentence that is said to explain this empirical success. However, attending to what counts as empirical success may introduce issues the Ramsey sentence approach is meant to avoid.

There are several general advantages to my account of scientific realism. However, the primary advantage is that it justifies the accumulation of scientific knowledge about the world and the entities within it, even if those entities are unobservable. This allows our scientists, engineers and technicians to have a justification for the knowledge they have about the world. From a philosophical perspective then, my thesis solves, or at least suggests a solution to, several problems relating to confirmation and approximate truth of scientific theories.

Chapter 1: Realism, Anti-realism and the Design of Instruments

1.7 Introduction

The central task of this chapter is to set out the problems to which my thesis is attending. Toward this end I will outline what I think is the most viable form of scientific anti-realism available – more or less the constructive empiricism offered by Bas van Fraassen.[1] An important aspect of this form of empiricism is its commitment to the existence of an external observable world. For many this must seem a welcome and obvious improvement on previous positivistic accounts in the philosophy of science. Yet many scientific realists (including this one) hope to go a little further in what we can believe is real on account of our best science. Like constructive empiricism, my realism emphasises the empirical basis of what we should believe as existing in the world. Though, as I will argue in later chapters, we can extend what counts as empirical to include the detectable world. This will require a defence on several fronts.

For a start, as well as the arguments in favour of constructive empiricism, there are also reasons why one should not be a scientific realist. One important kind of reason is based on what is termed the *pessimistic meta-induction* (PMI). The PMI challenges the realists' view of the history of science, where the success of a theory is explained by its approximate truth. The challenge to this view has led to the formation of several 'weaker' forms of realism. Yet some realists go so far in their attempts to accommodate the PMI that it is unclear how they might differ from anti-realism, or so I shall argue in later chapters. In this chapter, I will provide some preliminary comparisons that both clarify the nature of constructive empiricism and its relation to these 'weaker' forms of realism.

My working out of a realist position draws heavily upon Ian Hacking's version of entity realism. Hacking too has sought an empirical basis for what counts as real. His philosophy emphasises the importance of understanding the practices of engineers and technicians when resolving issues in the philosophy of science. Interestingly, the conclusion he draws from this analysis undermines the importance of theoretical concerns. According to Hacking, regardless of the truth of a theoretical description of an entity, if an engineer is able to *use* a hidden entity as a tool to further investigate the nature of the world, then the entity exists. In contrast, my thesis draws attention to the importance of the theories used in the design of instruments of detection and measurement. Furthermore, I argue that this theory has to be secure or fixed in order for us to be realists about what is detected and measured as a result of its proper use.

[1] I will draw upon a range of work but primarily as he develops this position in van Fraassen [1980].

So, although sympathetic to Hacking's position, ultimately I find that there are serious deficiencies with his account. Still, some of the concepts used in the development of entity realism are fruitful. To develop these ideas further, and thus put them to use in my version of scientific realism, I review some of Hacking's own influences. In particular, I briefly review the relevant work of the American Pragmatists John Dewey and Charles Peirce. The ideas presented here will be important in understanding the kind of 'instrumentalism' that forms an important part of my empirically based realism. Like Hacking I wish to emphasise the role of the physical instrument in detecting and measuring. Unlike Hacking I think we need 'secure' theory in order to understand what is indicated by the use of such instruments. Both Dewey and Peirce offer some important clues as to how this might be achieved.

1.8 Anti-Realism

In this section I provide some analysis of the kind of anti-realism that is the target of my thesis. What counts as an anti-realist position is an interesting and complex issue. I have chosen constructive empiricism over the other kinds because I am sympathetic to much of what it claims; however, I also think we can have more than what it provides. Like most anti-realist positions, constructive empiricism allows us to be realists about some things in the world. Roughly speaking, we are justified in believing that something exists if that thing is considered observable. As for the unobservable parts of the world, the constructive empiricist argues we should remain agnostic. It is important to notice that the agnosticism offered by van Fraassen does not rule out the existence of unobservable entities, it merely suggests that without empirical evidence we should not pass judgement on what exists in the empirical world. However, many realists, or at least this realist, do not want to remain agnostic about the existence of say, DNA molecules, just because they are unobservable. Observing an entity is a good reason to believe that it exists, yet this does not capture enough of what we ought to believe is real.

Before we examine constructive empiricism proper, it is worthwhile considering another important motivation for the anti-realist position – the pessimistic meta-induction. The response to the PMI offered by some realists, as we will see, sheds some light on the nature of anti-realism. Some realists have, in various ways, appealed to the success of science as evidence in favour of their position. It is contended that the success of various scientific theories would be a miracle if, in general, what these theories used to explain the phenomena turned out not to exist.[2] Yet as Larry Laudan [1981] has noticed, when we consider the history of science we discover many of the entities that are supposed to explain a theory's success turn out, upon further scientific investigation, not to exist. Some of these entities were thought to have been (at the time)

[2] This sort of miracle argument comes in many forms. Maybe there is no realist that would commit to this particular version. I use it here only to establish the dialectic. In chapter 5 I consider these matters in more detail.

well confirmed. Furthermore, although these entities were never discovered, they were important to the *success* of the scientific field that used them. Phlogiston, caloric, electromagnetic lines of force, various aethers, humors and miasmas are all good examples of postulated entities that were important (some more than others), but which, in the end, failed to materialise (in the light of our current science). It is clear from these examples that the existence of the postulated entity is not required to explain the success of a theory. The examples by themselves might motivate a form of scepticism. However, using the PMI the anti-realist can argue for something stronger: given that most past theories have been found to be false (and false to the extent that they can not be approximately true) we have good reason to believe that our current successful theories are likely to be false.

One realist response to the PMI is to agree with the conclusion but still note that empirical knowledge is grounded realistically. The unobservable can still explain the observable, though given our poor strike rate, we should not try and represent or describe the unobservable. This position has been termed 'modest surrealism'.[3] It is motivated by the idea that if we don't offer a representation of the unobservable it is hardly possible that we can be wrong about what exists. With this minimal claim regarding the unobservable the modest surrealist does get to explain, in some minimal sense, the success of science. All a modest surrealist need say is that *something* must exist to explain the observable phenomena, although they will abstain from saying anything further than this for fear of getting it wrong. A common way of accounting for this *something* is with the use of a Ramsey sentence – a formal device that separates theoretical commitments from empirical ones. I will consider a version of modest surrealism in chapter 5 and some problems with Ramsey sentence realism in general in chapter 6. Interestingly, the Ramsey sentence approach has also been used by some to try and represent the 'structure' of the unobservable world. This brings us to a second response to the PMI – structural realism.

There are several kinds of structural realism. For some we can commit to a structure that accounts for a similar *form* amongst different theories that range over the same empirical domain. For others, a structure can be inferred from the way that certain observational content is organised. At this stage we need not be too concerned about the relationship between the various modes of distilling structure. The important point to note is that, like the modest surrealist, structural realists are not (or purport not to be) committed to what it is that the theoretical terms in a theory refer to. Rather, there is some content regarding something real captured in a structure present in the various relevant theories. The relevance of modest surrealism and strict structural realism will become clear when we consider the agnosticism of constructive empiricism, to which I now turn.

[3] As far as I can tell the term 'surrealism' was introduced by Jarrett Leplin [1987]. Lyons [2002] develops this idea further to form the position he terms 'modest surrealism'.

Although an anti-realist, van Fraassen claims he is not an instrumentalist.[4] One of the reasons for dismissing this form of anti-realism seems to be that instrumentalists have to interpret theories in a non-literal way. Yet as van Fraassen himself notes when a scientist says ‘electrons are not planets’ or ‘there are electrons’ they really do want to be interpreted literally, rather than metaphorically or instrumentally [van Fraassen 1980: 11]. Even though theories literally construed do have a truth value, van Fraassen claims that:

> After deciding that the language of science must be literally understood, we can still say that there is no need to *believe* good theories to be true, nor believe *ipso facto* that the entities they postulate are real. [van Fraassen 1980: 11-12]

For van Fraassen, although theories can be true or false, the *acceptance* of them need not involve a belief in the truth of the unobservable part of the theory. This need not mean that truth does not enter into the assessment of a theory. Obviously what theories claim of the observable world has to be true – the theory has to be empirically adequate – but a constructive empiricist need only claim that;

> *Science aims to give us theories which are empirically adequate; and acceptance of a theory involves as belief only that it is empirically adequate* [van Fraassen 1980: 12] [italics original].

Empirical adequacy, although important for acceptance, is not usually sufficient. Apart from the belief that a theory is empirically adequate, acceptance, according to van Fraassen, also involves certain values. These pragmatic or instrumental values of acceptance relate to a theory’s predictive success, simplicity, explanatory power, scope, and so on. Importantly, acceptance also involves personal values, including a commitment, by the scientists involved, to the research programme. What is claimed by a theory in the empirical dimension has to be true and cannot contradict an observation. However, a scientist can be committed to the existence of a theoretical posit without empirical justification. Importantly, a commitment is not something that is true or false, it need only be ‘vindicated’ or not [van Fraassen 1980: 13]. A commitment to a theoretical entity is vindicated if it helps with the development of a novel prediction, unifies disparate theories, simplifies or explains and thus provides opportunity for further testing. In short, a commitment to the unobservable parts of the picture a theory draws is vindicated if that commitment helps increase the theory’s empirical adequacy. Thus, we have a way of believing what a theory claims to be observable, while measuring its acceptability as a total package. Although we might need to accept (and

[4] Maybe it is more accurate to say of van Fraassen’s anti-realism that he is not just an instrumentalist.

have reasons to accept) the theoretical posits beyond observation, we need not (for we have no justification) *believe* in their existence; we can and ought to remain agnostic.[5]

Still, it might seem that constructive empiricism is a form of modest surrealism in so far as it is agnostic about the unobservable cause of observable phenomena, while accepting that there is, at the very least, *something* behind the phenomena. I doubt, however, that this is what is being offered. Van Fraassen notes that realists have mounted several arguments suggesting that scientists, just by virtue of their scientific project, require a *common cause* to explain a statistical correlation of observables [van Fraassen 1980: 25-31]. Without recourse to further observables to explain the correlation, it looks as though science requires there to exist unobservable entities to act as this common cause. Van Fraassen counters this by suggesting that it is unlikely that the positing of an unobservable common cause will have no observable features. If this is the case, then we can recast the goal of science as developing models that might increase empirical adequacy. But now we have a choice between a common cause and an empirical adequacy account of science. To lend support toward the empirical adequacy account, van Fraassen argues that the indeterministic interpretation of quantum physics is a counter example to the common cause requirement.[6] Still, we might think that the common cause principle applies more generally. Yet van Fraassen has already established that, more generally, the common cause principle can be interpreted by the anti-realist as a pragmatic guide for theoretical model construction, or as an aid in the acquisition of empirical knowledge. The success of the former may depend on the latter, but neither justifies the belief in any unobservable metaphysical baggage.

It may be that van Fraassen is only making a point about the nature of explanation, and that his realism could extend beyond the observable world for other reasons. For example, some evidence might convince him to be a realist in a minimal sense about the unobservable world, but remain an agnostic when it came to any *particular* theory that explained the phenomena. The striking similarity in the equations used by different theories to account for the phenomena might suggest something common to these theories, even though each particular theory can be judged in terms of its empirical adequacy. As we have seen structural realists might take what is common to be some hidden structure. However, this too is a bridge too far for a constructive empiricist. Recently van Fraassen [2006, 2008] has developed what he terms ‘an empiricist structuralism’ to account for this similarity in the representations of the phenomena such that it has no realist implications. I will consider this approach in chapter 4.

[5] The concept of *acceptance* that van Fraassen is developing here is not without its problems [Musgrave 1985: 206-207, Maher 1990, Horwich 1991, Psillos 1999: ch. 9].
[6] In fact van Fraassen argues for something stronger. The probability calculus of the common cause account creates an inconsistency when applied to the indetermanistic interpretation of quantum physics.

Even though van Fraassen refers to himself as an agnostic, it now seems that there are several kinds of agnosticism. Maybe all forms of realism involve a degree of agnosticism. Even so, in order to make sense of constructive empiricism we have to be clear on the depth of agnosticism involved. I think it is helpful to categorise agnosticism into the following kinds:

(i) Moderate agnosticism,
(ii) Liberal agnosticism, and,
(iii) Strict agnosticism.

A moderate agnostic believes that because we have no epistemic access to the unobservable world, we will never know what the real or complete explanation is of the observable phenomena. Thus we should be agnostic about any particular posit used to explain the phenomena. Nevertheless, the moderates are prepared to admit that there has to be something unobservable to account for certain observable phenomena, and that these unobservables, just by virtue of the phenomena they produce, must have something in common. Structural realism and some forms of Ramsey sentence realism can be construed as attempts at moderate agnosticism.

For some, what passes as moderate agnosticism may be too strong. The fact that the moderates are prepared to classify a range of phenomena as having something in common just by virtue of an unobservable structure, looks as if they are making an epistemic claim about an unobservable. A moderate might claim there is an important relationship between the traces of ionising gas in a cloud chamber and the movement of an oil drop in an electric field. For the liberals, if this is just a pragmatic way of classifying the observations, all well and good. But if it is an epistemic claim about the unobservable world, then the moderates have gone too far. For all we know, the unobservable world is full of different entities responsible for what is observed in each experiment. To be sure, it might be convenient to classify these entities as similar on the basis of observable effects, but there are no epistemic grounds for supposing that the same kind of entity is involved in different experiments. We should not rule out *a priori* the scientific theory that suggests that each individual particle, field, molecule or whatever might have a unique set of properties accounting for each individual phenomenon. Liberal agnostics are modest surrealists in so far as they are committed to there being unobservable things accounting for the phenomena. However, they hold that we can infer nothing further about their nature.

The strict agnostic, like his religious analogue, goes further again. There might be no unobservable thing accounting for the phenomena. The strict agnostic notices that appearances are what appear to us, and that we cannot see beyond appearances.[7] It is

[7] This sentiment is expressed in van Fraassen's later work, for instance van Fraassen [2008: 99]. These matters will be considered in more detail in chapter 3.

true that perceptual judgements regarding an observable object may go beyond appearances, after all appearances can be deceptive. To this extent a strict agnostic is prepared to endorse simple perceptual judgements, at least in principle. However, we need not believe in any unobservable entities in order to account for what is observable. Constructive empiricism, it seems, is a version of strict agnosticism.

We can of course go further in our agnostic attitude so as to include the observable world.[8] However, van Fraassen rejects the notion that some sort of 'sense-data' theory might demarcate what counts as an observation. Van Fraassen is prepared to commit to some perceptual judgements regarding the observable world. Still there is a limit to the epistemological risks a constructive empiricist is prepared to take.

> There does remain the fact that even in endorsing a simple perceptual judgement, and certainly in accepting any theory as empirically adequate, I am sticking my neck out. There is no argument there for belief in the truth of the accepted theories, since it is not an epistemic principle that one might as well hang for a sheep as for a lamb. [van Fraassen 1980: 72]

Unfortunately van Fraassen says very little about what it is to make a simple perceptual judgement. In fact, he believes the analysis of what it is to observe is a question for science not philosophy. However, van Fraassen does give us some insight into the nature of observation as a constructive empiricist might construe it in his response to criticism from Grover Maxwell. Maxwell was actually targeting a positivist account of observation, but van Fraassen sees the criticism as relevant to the constructive empiricist.[9]

Maxwell's criticism turns on the 'observable-theoretical dichotomy'. If no clear distinction can be made between theory and observation, then the priority that the empiricist gives to observation collapses. Van Fraassen is clear when making precise just what is at stake here. Firstly, he notes that trying to establish a dichotomy between what is observable and what is theoretical is an example of a category mistake. Terms or concepts are theoretical, whereas objects or entities are observable (or not). What we have then is two distinctions; one between observation terms and theoretical terms, and another distinction between what is observable and what is unobservable. Of the first distinction van Fraassen agrees that observation terms are theory-laden, but nothing much follows from that, certainly not scientific realism. We might also note that even though observation terms are theory laden we should not confuse *observing* with *observing that*. If we were to show a tennis ball to some aboriginal people not at all

[8] Bertrand Russell [1927: ch. 20] argued that we should treat an observable object as a cause of our 'percepts'. He then considered other explanations for these percepts (for instance solipsism or phenomenalism). Thus, for Russell, what is observable is a matter of theoretical postulation.
[9] The original article from which van Fraassen gleans the relevant criticism is Maxwell [1962].

familiar with the game of tennis, it is fair to say that they would not see *that* it is a tennis ball. This does not mean that they do not see the tennis ball. Van Fraassen does not explicitly state the reason for this qualification, or indeed its relation to his theory; however, I think it is fair to assume that he is pointing to the 'theory-independence' of what counts as observable.

Maxwell [1962] also points out that the act of observation is rarely unmediated. We can look through an atmosphere, a window pane, spectacles, a magnifying glass, binoculars, a telescope, a low-powered microscope, right through to a high-powered microscope. The question then is, is there a principled way of determining when this 'looking through' is so mediated that we can no longer classify what is observed as the observable entity? Van Fraassen's response to this is to argue that Maxwell has shown that observability is a vague predicate. As long as we can give clear cases of what is observable nothing much follows from this. Examples are duly provided: An electron is unobservable for it is in principle not possible to observe it, while the moons of Jupiter are observable because if we were close enough we could observe them.

The second argument van Fraassen identifies in Maxwell's paper turns on the fact that what makes something observable is context dependent. Many things, for example, are not observable in a dark environment, or are observable only under special lighting conditions. The current technology of telescopes also limits what is observable in our universe. So what if we were to evolve electron microscopic eyes? Would this not make more of the world observable? Van Fraassen thinks this is just a trick. We don't call the Empire State Building portable just because we might evolve into giants. The 'able' in observ*able* is a human ability in so far as we understand what it is to be a human. Yet there does seem to be a bit of wriggle room here. It is currently impossible for us humans to observe the moons of Jupiter (the technology does not exist for us to get close enough to the moons) but if and when it does arrive, we are allowed to use it to ensure that those moons are in fact observable. So this (human) ability of observation can, to some extent, depend upon certain technologies. In chapter 3 I intend to draw upon this use of technology to extend our abilities, especially to those detectable parts of the world.

Even though van Fraassen has given us some clear examples of what is observable and what isn't, he says that it is in fact up to the relevant epistemic (read scientific) community to decide on this matter. Given that, as van Fraassen admits, all observation statements are theory-laden, and that there is no clear distinction in scientific language between unobservable and observable, it looks like we have to either just stipulate what counts as observable or find some way of escaping what seems like a vicious circle. Van Fraassen notes that:

> To find the limits of what is observable in the world described by theory *T* we must inquire into *T* itself, and the theories used as auxiliaries in the testing and application of *T*. [van Fraassen 1980: 57]

If a theory *T* describes the limit of what is observable, one might wonder how independent this description is from theory *T*. The appearance of circularity here is only apparent. What is observable is just a matter of the *facts* about us, and these facts are 'disclosed' by the theory itself. Apparently the disclosure of these facts is a theory-independent question. What this independence amounts to, it seems, is the point being made in the case of the aboriginal observing a tennis ball. Although suggestive, the relationship between the examples van Fraassen points to and the epistemic issues they are supposed to solve is under developed.

Notwithstanding van Fraassen's emphasis on science itself sorting out what counts as observable, it is clear that perceptual judgements and their relation to what is observable are important for the constructive empiricist. Part of the project of the thesis will be to fill out, in a bit more detail, a viable account of perception consistent with the little that van Fraassen has said on the subject. I will then argue that analogous features can be found in a viable account of detection. If the analogy is sound, a constructive empiricist could recast his theory in terms of detectability so as to remain agnostic about what is not detectable. Maybe this is asking van Fraassen to stick his neck out too far. Still I think it is consistent with what scientists take an 'observation' to be, and if I am successful in arguing that detection is analogous to observation, the onus is back with the constructive empiricist to provide relevant reasons why we should maintain the distinction between the relevantly similar kinds of activity. Key, then, will be the development of a model of detection suitable for such an analogy. Some of the insights provided by Ian Hacking's entity realism will form a central part of the story.

1.9 Entity Realism

The power of the pessimistic meta-induction lies in the assertion that most successful theories have turned out to be false. This, on the face of it, undermines those types of realism that point out that the success of science has something to do with the truth (or more accurately approximate truth) of the theories involved. Hacking's development of entity realism can concede that a theory can turn out to be false, yet this need not affect our ability to genuinely refer to the real and unobservable entities of the world. To appreciate this point, what we should focus on are the practices of experimental scientists. Experimentalists, apparently, often use unobservable entities as tools to investigate further the hidden nature of the world. For these experimentalists, there is no important problem with having a theoretical description of an entity turning out to be false. For them, the ability to use an unobservable entity as a tool provides enough warrant for the assertion that it exists. Even though I may harbor an incorrect theory regarding how a thermometer works, I can still use it to measure a temperature, and as a result I can commit to its (the thermometer's) existence. So it goes for many unobservable entities also used as tools. This understanding through 'intervening', as Hacking terms it, was learnt from Dewey. A way of immunising this knowledge from changes in theoretical description was provided by Hilary Putnam.

If theoretical descriptions of an entity determine what we mean, then this allows what Hacking terms 'meaning incommensurability' to take hold. If our theoretical terms are implicitly defined by the theories in which they are embedded, then two scientists who hold different theories mean something different by the term 'electron'. Different scientists, just by virtue of their particular theory, are describing different things in the word. If this is right, with each new development in a theory we seem to be committed to there being different things in the universe. However, as Hacking notes, although J.J. Thomson did not use the term 'electron' in his theory, and although there have been many theories about electrons,

> the incommensurabilist would be out of his mind if he said that Thomson measured the mass of something other than the electron – our electron, Millikan's electron, Bohr's electron. [Hacking 1983: 84]

In contrast, Hacking notes that Hilary Putnam has provided a credible way of understanding how we might refer to the same entity, even if there are varying descriptions amongst language users. According to Putnam, the meaning of a word or term has four components. Firstly, it has a *syntactic marker* that classifies the word grammatically. Secondly, a *semantic marker* describes just what kind of thing is being referred to. This can be quite general, as in a natural kind, or a more specific kind, such as the kind 'liquid'. The third component is the word's *stereotype.* This is a convenient, though potentially inaccurate, characterisation that language users can use when referring to what the word means. The stereotype might note common, though not necessarily accurate, attributes of what is being referred to, such as 'a tiger has stripes'. It is clear that Hacking thinks of a theoretical description of an entity as its stereotype. Like all stereotypes they can be developed, even revoked or reformed in light of what is discovered about them. The fourth component is that to which the term refers. It is the method of referring that is doing the real work in immunising us against meaning incommensurability. Reference requires a connection or causal chain of the right sort between what is being referred to and the language user. A cloud, for example, is free to cause certain effects (observations) that we can describe. Whether two individuals, describing a cloud, are in fact referring to the same cloud, depends on the sort of (causal) connection between them and the cloud(s) in question. If they are both referring to the same cloud, they might disagree on the description of it, but not what they are referring to. Reference is of course more complicated than this, especially when we consider that much of what we are referring to is established by others. Even so, Putnam's solution seems a reasonable 'base case' explanation of how we might refer to an object.

Although this solution can work reasonably well at the level of the observable, scientific terms often refer to *unobservable* entities. Given this, it is not clear how the above relationship between the observed and its description might work for the unobservable posits of science. Influenced by the early American pragmatists (Dewey in particular), Hacking proposes that a causal relation can be established, even if we cannot *observe* what we are trying to refer to. Rather than observe, we can achieve a similar relationship by *doing* things with the relevant unobservable entities. In particular, if we can use them as tools when building experimental technologies, then reference to the entity can be achieved.

Part of the power of Hacking's argument is the detailed descriptions he provides of how complex instruments are developed by engineers and technicians. From an outside perspective it is difficult to understand how they might achieve what they achieve without there being unobservable entities that allow the technology to be built and function. From the engineer's perspective there is no inference involved; they are just *using* the hidden entities of the world if in fact they can. Importantly, according to Hacking, the various technicians and engineers involved in the construction and use of a piece of equipment need not have a common set of beliefs about what they are using. Hacking's favourite example is the *use* of an electron in various experimental technologies. Apparently we could ask of these experimenters;

> [M]ight there not be a common core of theory, the intersection of everybody in the group, which is the theory of the electron to which all the experimenters are realistically committed? I would say common lore, not common core. There are a lot of theories, models, approximations, pictures, formalisms, methods and so forth involving electrons, but there is no reason to suppose that the intersection of these is a theory at all. [Hacking 1983: 264]

Given the lack of theoretical overlap, it is unlikely that it is a description that establishes for these engineers and technicians that they are referring to the same kind of entity. Nor is it the truth of any description that justifies the belief in the existence of the entity. Rather, it is our ability to manipulate the electron that establishes for us its existence. Hacking's appreciation of this point came via his inquiries into the nature of a particular experiment involving charged particles;

> How does one alter the charge on the niobium ball? 'Well at that stage' said my friend, 'we spray it with positrons to increase the charge or electrons to decrease the charge.' From that day forth I've

> been a scientific realist. *So far as I am concerned, if you can spray them then they are real*. [Hacking 1983: 23]

If Hacking is right, our experimental practice ensures the existence of the entities that are being used, while also providing the right causal connection in order for us to genuinely refer to the relevant entity. To appreciate this way of establishing the existence of a scientific entity is to be an *entity realist*. On the other hand to accept the truth (or approximate truth) of the theories that describe these entities is to make a commitment to *theoretical realism*. Having separated the two kinds of realism, Hacking contends that, regardless of the truth or falsehood of the theory that describes something, if you can use something as a tool you can commit to its existence. If this is right, it looks as though we have a way of committing to entities that can survive the most radical change in theoretical description. Moreover, given that entity realism is not committed to the truth of a theory, it looks as though it can avoid the anti-realist implications of the PMI.

Unfortunately, the above solution might be both too weak and too strong. If there is no relationship between theory and entity, it is not clear that we can know *what* we are referring to. We can, I suppose, say that, by being able to use something we are referring to something, but this is to say very little indeed. This limited form of realism we might term 'pure entity realism'. If all a pure entity realist claims is that, when we are experimenting we are doing something with something, then only a very weak form of realism is being endorsed. A liberal agnostic might agree with this much realism. Unfortunately, the weakness of pure entity realism also makes it too strong. It allows anyone who is *using* an entity (by spraying or whatever) to have a referent for their theory. For instance, an experimentalist might claim that caloric can be added to materials, and thus used to make it expand. Some theorist thought that caloric could be used to determine other properties such as the relationship between latent and specific heat in certain materials. (A version of the caloric theory explained quite successfully the velocity of sound in air.) The pure entity realist is forced to understand the term caloric as a name that refers to heat as a form of energy (the real entity being used). Yet caloric theory is now considered false because it does not genuinely refer to the type of thing that was meant to explain the phenomena. If, contrary to this assessment, pure entity realism allows the term caloric to refer to whatever it is that is actually used by the experimentalist then it is far too permissive.[10] The problem can be seen from another perspective. Unless we describe an entity as having a set of specific causal properties, it is not clear how we might assess existence claims regarding the entity, especially if such assessments have nothing to do with the truth or falsity of such descriptions.[11]

[10] Laudan [1984: 156-162] makes a similar criticism of Hardin and Rosenberg's [1982] equally permissive use of a causal theory of reference.

[11] For related and further criticism see Resnik [1994].

Hacking's interpretation of the caloric situation is perplexing. He argues that because the theorists using the term caloric were *not* referring to anything, Putnam's theory of reference does not work. In these sorts of cases we can revert to the 'language game' (read descriptive) account of meaning [Hacking 1983: 87]. But now it is not clear how Putnam's account is supposed to help. How are we supposed to decide whether the right sort of causal chain is operative, or, alternatively, whether we are just playing a language game (while playing with something else in our experiments)?

In other places, Hacking seems to strengthen the link between entity realism and our description of the relevant entity:

> *We are completely convinced of the reality of electrons when we regularly set about to build – and often enough succeed in building – new kinds of devices that use various well-understood causal properties of electrons to interfere in other more hypothetical parts of nature*. [Italics original] [Hacking 1983: 265]

It is not clear how we might understand these causal properties without a theory, or some propositional knowledge, to interpret them. In fact if Hacking is right, engineers do use knowledge of causal properties in their design of devices. In other places Hacking does seem to recognise this, conceding that we need a 'modest number of home truths' in order to design and build [Hacking 1983: 265]. I claim that there is a lot more to these 'home truths' than that to which Hacking gives credit. I think there are good reasons to believe that this theory is not just 'lore' either. It takes quite a lot of sophisticated theory to measure the mass of 'our electron', especially to the accuracy that our modern instruments provide.

How might we strengthen this relationship between the entity and the theory describing it? Although we might worry about theory change in the pure sciences, there may be a method that ensures that the theory used in a design is 'secure'. There is a *prima facie* case for this idea given that there are many design principles that have survived scientific revolution. For example, many of the laws of classical physics are retained for designing in the classical world. For that matter there is much design knowledge available for designing instruments that detect and measure in the non-classical world as well. If this is right there is no reason to suppose that we can't detect and measure certain properties of the unobservable entities of the world using such theory. Much work will be required to understand what it is to 'secure' a theory such that we can detect or measure unobservable properties of these entities and the causal relations they enter into. I will also need to be careful not to beg the question against the constructive empiricist. Fortunately there is a form of instrumentalism that allows a limited amount of realism that might help secure or fix a theory without begging any important questions.

In the next section we will see how trying to provide an account of securing or fixing a theory led Charles Peirce to a form of realism that betrays commonsense. The development of John Dewey's instrumentalism owes much to Peirce's pragmatism, but it is also a response to its problems. Like van Fraassen, Dewey emphasises the pragmatic virtues of a theory. Moreover, he too draws attention to the important role of the measuring *instrument* in science. However, by drawing upon some of the insights of operationalism and his focus on the applied sciences, Dewey has provided a way forward for an empirical basis for scientific realism.

1.4 The Realism of Pragmatism

Although primarily considered a philosopher, Charles Peirce was also involved in some important science. As a result, Peirce appreciated that scientific enquiry had to manage experimental error and varying interpretations of the phenomena involved. Even so there were important aspects to the scientific method that could be utilised in a philosophical project. Broadly speaking, for Peirce's pragmatism, all beliefs we form about the empirical world must be the consequence of a real world investigation. A second and related project was to justify the use of the pragmatic method itself. Unfortunately, a commonsense realism that might have provided the foundation for this method gave way to an obscure and complex ontology.[12]

The most important statement of Peirce's pragmatism comes in the form of two papers, 'The Fixation of Belief' and 'How to Make Our Ideas Clear', first presented as a single essay in 1872 at The Metaphysical Club.[13] According to Peirce, the name of the club was intended ironically, with all attendees sceptical of any foundation for metaphysics. In the paper, 'The Fixation of Belief' Peirce considers four ways that might effectively eliminate real doubt and thus fix a belief. The first method – the method of tenacity – relies on the single-minded dogmatism of individuals to pursue beliefs despite other people's doubts. Although this is effective to some extent, it does depend upon a degree of intellectual isolation. Peirce believed that, in practice, this method was impossible to maintain. Just by virtue of the exposure to other people's different ideas, our tenacious hold on a set of beliefs will be undermined. Another possibility was the method of authority. Here, we might have some external authority suppress doubt and thus provide consensus through sanctions and rewards. In practice, history has taught us that, like its tenacious counterpart, it is impossible to maintain. The third method is to allow absolute intellectual freedom. Any belief is justified if it is 'agreeable to reason', where minimal standards such as consistency are the only constraint. When forming beliefs about the

[12] It is in fact unclear what Peirce's realism amounted to. For a thoroughly realist interpretation of Peirce, see Legg [1999]. For a more idealist, or maybe Kantian interpretation of Peirce see Hookway [1985].
[13] First published in *Popular Science Monthly* in the late 1870s. References are to their republication in Peirce [1998].

world this method is far too permissive. In fact it does not seem capable of actually *fixing* any belief.

From a pragmatic perspective then, none of these first three methods is effective in alleviating real doubt and thus fixing a belief. It is clear that to do the job we need something independent of the individual involved.

> To satisfy our doubts, therefore, it is necessary that a method should be found by which our beliefs may be caused by nothing human, but by some *external permanency* – by something upon which our thinking has no effect [my italics] [Peirce 1998: 25].

A little further on, Peirce goes on to claim:

> Our external permanency would not be external, in our sense, if it was restricted in its influence to one individual [Peirce 1998: 25].

The fourth and final method is the scientific method. Importantly, this method allows doubt to feature in the form of a hypothesis, while providing operations for the resolution of this doubt. Also required, as suggested above, is an 'external permanency'. These two aspects then form the basis of the pragmatic principle. It is an external permanency that when coupled with the scientific method enables us to fix an empirical belief. Although what constitutes the scientific method is largely a matter of investigation, Peirce, especially in his later writing, is rather speculative about the nature of the external permanency and the mechanism by which it helps fix a belief:

> The real, then, is that which, sooner or later, information and reasoning will finally result in, and which is therefore independent of the vagaries of me and you. Thus the very origin of the conception of reality shows that this conception essentially involves the notion of a COMMUNITY, without definitive limits, and capable of definite increase of knowledge. [Peirce 1960: 186-187]

Peirce might only be making an epistemic point, but then a little further on:

> Finally, as what anything really is, is what it may finally come to be known to be in the ideal state of complete information, so that reality depends on the ultimate decision of the community; so that thought is what it is, only by virtue of its addressing a future thought which is in its value as thought identical with it, though more developed. In this way, the existence of a thought now depends on what is meant to be

> hereafter; so that it has only a potential existence, dependent on the future thought of the community. [Peirce 1960: 189]

Given that what it is to be real now depends on an 'ideal state of complete information', it looks as if Peirce is offering us a form of idealism. Even more problematic is the notion that it does not exist now but in the future. This evolving conception of reality, Peirce terms *agape*. There seems to be two ways of understanding this notion. In so far as our beliefs evolve through scientific investigation, the external permanency that depends on the 'ultimate decision of the community' is also the result of an evolutionary process. Alternatively, the external permanency is itself evolving, with our co-evolving beliefs fixed through the scientific investigation of the external permanency.

This speculative philosophy seems the exact kind of metaphysics that the common sense pragmatic principle was meant to guard against. What then is the source of this strange ontology? It seems that, like the positivist's reliance on the verification principle to provide meaning for *all* statements, Peirce's ontology was a result of thinking that the pragmatic principle (external permanence coupled with the scientific method) could fix all empirical beliefs. Although related to testing of what exists, some beliefs – those beliefs about the past, or the future, dispositions or what is physically possible – require more than the pragmatic principle can provide.

The beginnings of the problem are evident in his paper 'How to Make Our Ideas Clear'. Towards the end of this paper Peirce considers the case of a diamond that exists for a time and then is destroyed without its hardness having been tested [Peirce 1998: 46]. According to the pragmatic principle, to believe the proposition 'this diamond is hard' is to have tested the diamond for its hardness. It is the external permanence (the test with an agreed-to interpretation) that justifies (fixes) the belief that 'this diamond is hard'. Unfortunately we can still ask whether an untested diamond is hard or soft. What property are we referring to, when we say of a diamond, which has not been tested for hardness, that 'this diamond is hard'? To answer this question Peirce offers a thought experiment. Let us say that every thing in the world is in fact soft; however, some things when tested immediately become hard. The pragmatic implication is that, if there is no testable difference there is not a real difference. Hence, it makes no difference to say that an untested diamond is soft!

Ultimately Peirce is unsatisfied with this conclusion. He does in fact want untested diamonds to be hard. In a later paper Peirce suggests that we are justified in considering an untested diamond as hard if it falls under what he terms a *general* [sic]. In the case under consideration the general would be, all diamonds are hard [Peirce 1960: 305ff.]. Yet this serves only to sharpen the problem. We are now in need of an external permanency to fix the belief in this general. Moreover, the general has to cover all time and so applies to those diamonds yet to be created. How can an *external permanence* justify a general, when a general encompasses both temporal facts about the properties of future diamonds, as well as unrealised properties – causal and dispositional properties – of actual diamonds?

Peirce's solution to this problem involved, what he termed, *synechism* [Peirce 1955]. This approach is best understood via an analogy with mathematics. If we consider a collection of points, the gaps between the points can never be filled in by adding more points to make a line. On the other hand, we can 'fill in' these gaps with a mathematical function that *generates* a curved or straight line to fit the points. Something similar is happening with abduction from data points to a law of nature, or facts to a general. However, to fix the belief in a general, we need more than abduction, we need an external permanency. There are after all many lines that can be abducted consistent with the data. Synechism emphasises that in the same way that the truth maker of a mathematical function describing a line is not a series of points, the truth maker of an abducted hypothesis is not the set of data forming the basis of the abduction. If a general correctly refers, we have a referent to future facts (predictions), as well as having a referent and justification for dispositional properties. As more data are fixed by the external permanence there will be some natural variability, yet the community is allowed to adjust the general so that it is enriched and thus evolves towards the truth.[14]

This analysis of the nature of the external permanency seems to raise more problems than it solves. In later writings, Peirce considers the external permanency as a sort of *dynamical object* that, according to some, is more like a process than an immediate object.[15] Each individual diamond, for example, is a process or continuum through time, that collectively, like some Gaussian surface (to carry the mathematical analogy), makes up a dynamical object that acts as the truth-maker for diamonds in general. These matters are complicated further as Peirce suggests that the community is able to enrich the dynamical object itself through ongoing testing of diamonds and the beliefs that are fixed as a result. Here then, we have something like evolutionary idealism.[16]

This synopsis does not do justice to the subtlety, scope and complexity of Peirce's work. Even so, it is clear enough that Peirce's form of realism is problematic. As some of the problems Peirce grappled with will test my own theory, I am wary of inflating my ontology beyond the common sense that realism is supposed to use as its touchstone. Like Peirce, I wish to use a pragmatic principle to secure or fix a theory (if indeed we can) when regarding something real. Thus, I agree with Peirce that an external permanency, once detected, can fix or secure a belief. However, in so far as we can infer that something is observable in virtue of our observations, and that this is enough to

[14] As mentioned above there is a double and maybe even reciprocal evolution going on in Peirce's ontology. For a more fine grained analysis see Hausman [1993].
[15] For a further development of these ideas and their complexities see Hausman [2002: 13-27].
[16] For an interpretation that does not attempt a reconciliation of these complexities see Miggotti [1998].

commit to its reality, then a similar inference, at least in principle, is available regarding what is detectable. Whatever else is required we can leave to the metaphysicians. Although I think this response is appropriate when attending to the concerns of the constructive empiricist, in order to distinguish differential realism from other similar forms of realism, I will also provide some analysis of its ontological commitments. These issues will be the central concern of chapter 4.

1.5 The Instrumentalist Fix

Dewey, a student of Peirce, reacted against Peirce's ontology while maintaining the pragmatic core. One could avoid the speculative metaphysics if one denied that the pragmatic principles provided any insight into the true nature of the world. Indeed Dewey thought that the concept of truth was more a hindrance than a help to any inquiry [1981: 31]. Given the somewhat chaotic nature of the world we experience, the best we can hope for is *warranted assertibility*. Some belief is warranted if it helps solve some problem or is useful as an *instrument* of inquiry. We can create these instruments in the form of theories or explanations, or more importantly, in the form of a technology or technique. For Dewey, trying to assess these instruments beyond their ability to solve a problem is to misunderstand both the real value of these instruments and the nature of inquiry itself.

Hacking, influenced by Dewey's approach, also notes that there is a bias towards 'knowledge as representation of the world' in philosophy [Hacking 1983: 130-1]. This bias, according to Hacking, Dewey parodied as the 'The spectator theory of knowledge'.[17] Rather than hope that some representation might correspond to some reality, Dewey wanted to emphasise a 'genuine correspondence', an intimate knowledge one has by virtue of negotiating, manipulating or controlling an environment. Here the world and self really do *co-respond* in a functional or transactional way. Dewey's ultimate concern is the development of a way of knowing through doing – a form of know-how he thought characteristic of the applied sciences. As we investigate nature we are faced with practical problems. In response we formulate ways of overcoming, controlling and thus understanding these problems. It is only through manipulating with the help of tools, and solving practical problems, that we begin to understand what we are investigating. Dewey's instrumentalism, then, really does involve physical instruments.

Dewey well understood that for a science to be effective not everything could be questioned at once. In order to investigate scientific problems some assumptions have to be fixed. To control for these assumptions, scientists need only determine them

[17] Hacking goes on to state his debt to Dewey: 'My own view, that realism is more a matter of intervention in the world, than of representing it in words and thought, surely owes much to Dewey' [Hacking 1938: 62].

operationally. After suggesting that we gain commonsense knowledge just in virtue of manipulating objects, Dewey notes:

> The important thing in the history of modern knowing is the reinforcement of these active doings by means of instruments, appliances and apparatus devised for the purpose of disclosing relations not otherwise apparent… [Dewey 1987: 87]

Dewey goes on to claim that:

> Among these operations should be included, of course, those which give a *permanent register* of what is observed and the instrumentalities of exact measurement by means of which changes are correlated with one another [Dewey 1987: 87; my italics].

A guiding principle in Dewey's work seems to be that we can map onto a physical instrument an operation that determines the meaning of a concept. This establishes a relationship between a concept as an idea and its concrete manifestation. According to Dewey '[a] definition of the nature of ideas in terms of operations to be performed and a test of the validity of the ideas by the consequences of these operations establishes connectivity within concrete experience' [1987: 114]. In fact, '[t]heoretical certitude is assimilated to practical certainty; to *security*, trustworthiness of instrumental operations' [Dewey 1987: 128].[18]

Although it is not always clear what this correspondence between theory and instrument amounted to, it does seem that he thought there was a physical structure to our world. We could take advantage of this structure to provide warrant for the predictions of a theoretical system. Through the development of new technologies we can fine-tune aspects of our theories to fit the ever emerging detail of science. In so far as we can rely on this structure of the world, and upon which we can map a theory, Dewey was a realist.[19] However, this process was not guided by, or getting closer to, some truth or ultimate knowledge; rather, the process was allowing us to get better at solving problems both in detail and scope.

Although I think Dewey is on the right track, I hope to show that there needs to be more to operationalism than just a mapping of a theory onto an instrument or physical operation. The kind of operationalism Dewey had in mind, after all, has suffered considerable criticism. Yet there is a way of understanding operationalism based on what Schlesinger [1958] terms the 'differential approach' that survives much of this

[18] For further development of what this mapping of concepts onto physical operations meant for Dewey, see Tristan [2002].

[19] Dewey is usually interpreted as an anti-realist. According to Godfrey-Smith [2002], Dewey emphasis on relations could be construed as a form of structural realism.

criticism. When we couple this approach with Dewey's insight regarding the role of the physical instrument or physical operation, we are provided with a way of securing a theory. This secure theory justifies what we interpret as indicated by the use of an instrument of detection or measurement.

Instrumentalism and Realism

For some, my appeal to the instrumentalism of Dewey might seem anathema to the realist cause. Dewey's emphasis on the instrumental value of a theory might seem to undermine the view that a theory has some important truth-assessable relation to the unobservable features of the world. The instrumentalism that Dewey developed considered it the case that the effective application of a theory was more to do with know-how, and know-how is not the kind of thing that can be true or false. This view, however, need not affect the justification for the form of realism I hope to develop in the next two chapters. The realism I am developing begins by extending what a constructive empiricist can consider as 'observable' and thus real. If this sort of instrumentalism is important when considering the observable world then the same approach need not undermine, indeed it may support, judgements regarding the detectable world.[20]

Still there are important issues at stake when taking an instrumental view of a theory. One might wonder, as Popper did, what it is to test a theory if it is viewed as an instrument. In the next chapter I will argue that Popper was mistaken in thinking there was a dichotomy between instrumental and thus unfalsifiable know-how on the one hand, and risky propositional knowledge awaiting a severe test, on the other. There is a third option: there are testable theories that, because of such tests, can be evaluated for their applicability. I will argue that testing an instrument can provide us with important information regarding its applicability, and that this is a genuine form of testing distinct from what Popper construes as a severe test

[20] Arthur Fine argues that 'at heart constructive empiricism is a version of Dewey's instrumentalism' [Fine 2001: 112].

Chapter 2: The Test of an Instrument

2.1 Introduction

The purpose of this chapter is to defend the proposition that we have *secure* theory describing the empirical world. Secure theory will be important when, in the next chapter, I defend the proposal that what is detectable as well as what is observable constitutes the 'empirical world'. To establish this I require fixed or secure theory that describes how our instruments of detection and certain physical processes function. This extension of *the empirical* means that a theory that is empirically adequate can inform us about, not only what is observable, but also what is detectable. However, if the description of what is detected is problematic in a way that the description of what is observed is not, then the constructive empiricist has reason to remain sceptical about the contents of the detectable world. There is, though, another type of sceptic who might scupper the whole project. The falsificationist will argue that there is no secure knowledge based on observation or otherwise.[1] If successful, falsificationism would derail both constructive empiricism (usually construed) and my extension of it. So that I don't beg the question against the constructive empiricist while I defend against the falsificationist, I will only argue that we have fixed or secure theory in so far as it is relevant to the empirical world. I leave open, in this chapter, what it means to be empirical.

The sceptic of this form of secure theory I take to be Karl Popper. Popper acknowledges that a large body of supplementary theory is required to test a theory, and although he characterises this kind of knowledge as 'unproblematic background knowledge', by this he only means we should *accept*, tentatively, this knowledge as unproblematic. But even this tentative acceptance is risky: firstly, because it is not possible to confirm a theory; and secondly, because any attempt to accept a theory as established would hinder the development of science. I think Popper is probably right about the first point, but this does not mean that we can't be justified in believing certain aspects of a theory are correct. Moreover, this is important for scientific development.

The central goal of the chapter then, is to show that Popper's falsificationist programme requires the kind of secure theory that I am advancing. If this is right the scepticism of falsificationism is truncated so as to consider just those contexts where the application of a theory may be problematic, while leaving aside established theory. There are three related arguments for this proposal.

[1]The incommensurablist may also oppose this thesis [Kuhn 1970, Feyerabend 1975]. I will not provide an argument against this view of scientific development. Having said this, I think many of the points raised would disarm the concerns of the incommensurability thesis.

The first argument (§2.3) concerns the role of the 'severe test' in Popper's account of science. I shall argue that without a different kind of test, or a different interpretation of the results of a severe test, we cannot explain the unique success of the applied sciences. I term this second kind of test an *informative test*. Roughly, the informative test provides information about the applicability of an instrument, where an instrument can include a certain type of theory or model. This latter claim is not meant to suggest that a theory or model is a physical instrument. However, this form of instrumentalism does note an important relationship between a theory, construed instrumentally, and its use in relation to a physical instrument. What is involved here will be dealt with in more detail in the following chapters. Here, we need only note that according to Popper's approach, a severe test is the only genuine test we have. Moreover, it can falsify (or corroborate) a theory or establish its applicability. Unfortunately, Popper's approach makes applied science either a non-science (because it doesn't severely test) or indistinguishable from pure science (because it does). Either alternative, I shall argue, is unacceptable.

The second argument (§2.4) considers Popper's ruling against the *ad hoc* development of a theory. I show that this rule conflicts with another important aspect of his programme: what it is to 'easily test' a theory. The informative test identified in the first argument accounts for the easy test in a way that avoids this conflict. The third argument (§2.5) considers Popper's own account of how a falsificationist should understand the reproducibility of a test. It turns out that the reproducibility of a test also depends on the informative test. The falsificationist may have a response to these last two arguments, but only if they can answer the dilemma raised in the first argument.

Even if the argument is successful, it may still seem a mystery how we assess what might count as a secure theory. Here I briefly consider the work of Deborah Mayo, who, utilising many of Popper's insights, has provided a way of measuring the degree of support a hypothesis garners from a test. Although sympathetic to her project, I argue we need more than Mayo provides. In the next chapter I develop a form of operationalism that provides a way of testing, and thus securing, important parts of a theory.

2.2 Popper and the Demarcation Principle

Popper thought that distinguishing science from non-science was a matter of providing an account of the rules or methods used in scientific practice. A demarcation principle, then, identified something science did and non-science failed to do, the result of which made the latter less, and the former more, scientific. Pivotal to Popper's account is his understanding of the development and role of a 'severe test'. We will consider shortly what such a test involves. For the moment, it is enough to note that Popper believed that it was severe testing that was important in demarcating scientific practice from non-science.

Popper also contrasted his demarcation principles with the programme developed by the logical positivists. Unlike the logical positivists, Popper suggested that we ought to consider *all* scientific statements as unjustified conjectures:

> We shall have to get accustomed to the idea that we must *not* look upon science as a 'body of knowledge' but rather, a system of hypotheses; that is to say a system of guesses or anticipations which *in principle cannot be justified* … [my italics] [Popper 1965: 317]

Again:

> Predictions… based upon universal theories and initial conditions are not 'justified', or 'warranted', in any sense of these terms, any more than universal theories themselves are … [J]ustification need not play any role in the critical analysis of scientific knowledge. [Popper 1979: 364]

What seems like a weakness Popper sees as a strength. Consider the universal statement 'all massive bodies are subject to gravitational forces'. It is difficult, if not impossible, to confirm this statement when considering all space and time. Even so, Popper points out an asymmetry here: although 'all massive bodies are subject to gravitational forces' can never be confirmed, we may be able to find conditions under which it might be false. Clearly, if we were to find a massive body unaffected by a gravitational field, then we could falsify the above universal claim. The first stage of the scientific method involves finding conditions under which, and within their scope, theories are falsifiable. Theories can be considered scientific in so far as there are conditions under which it might be found to be false. A theory can be falsified by testing under just these conditions.

This rather straightforward characterisation of the demarcation principle requires further qualification. For a start, a statement that might be used to falsify our theory, for example, 'this massive body is unaffected by gravity', involves theoretical concepts. Even a very simple statement like 'this is a glass of water' is theory laden [Popper 1963: 119]. What makes something massive, or glass, or water, involves a large range of dispositional properties, and the description of a dispositional property, according to Popper, is theoretical if anything is. Given that we can never confirm a theory, it now looks as though we can't confirm any falsifying statement. Testing a theory then, can only ever be an *attempt* to prove it wrong. Falsification is not about determining that a theory is false, rather, it is about providing the conditions whereby we can test the theory. If it fails the test, we have reasons to accept (or decide) that the theory under test is falsified relative to the accepted, though problematic, assumptions of the test.

Although all knowledge is problematic we can't question it all in any one test. In order to get on with the business of science we have to partition off that knowledge or part of our theoretical system we want to test, from what we accept as background theory. Although accepting problematic background theory as unproblematic seems risky, this does not worry a falsificationist.

> The fact that, as a rule, we are at any given moment taking a vast amount of traditional knowledge for granted (for almost all knowledge is traditional) creates no difficulty for the falsificationist or fallibilist. For he does not accept this background knowledge; neither as established nor as fairly certain, nor yet as probable. He knows that even its tentative acceptance is risky, and stresses that every bit of it is open to criticism, even though only in a piecemeal way … [Popper 1963: 238]

If this is right there is a huge range of things that can go wrong with an experiment. With all knowledge equally problematic, a scientist faces a significant problem when trying to determine what has been falsified. Yet Popper points out that deciding how a theory performed in a test is made easier in several ways. There are ways of increasing the degree of falsifiability, and thus testability, of the theories involved. Popper uses an astronomical example to articulate this idea. Consider the following four statements construed as laws of nature:

p: All *orbits of heavenly bodies are circles*.[2]
q: All *orbits of planets are circles*.
r: All *orbits of heavenly bodies are ellipses*.
s: All *orbits of planets are ellipses*.

According to Popper:

> Moving from *p* to *q* the degree of universality decreases; and *q* says less than *p* because the orbits of the planets form a proper subclass of the orbits of the heavenly bodies. Consequently *p* is more easily falsified than *q* … Moving from *p* to *r* the degree of precision (of the predicate) decreases: circles are a proper subclass of ellipses. [Popper 1965: 122]

The lesson seems to be this: we can increase the falsifiability of a statement by expanding its scope from planets to all heavenly bodies. This increases the range or scope of potential falsifiers. But we can also increase the falsifiability of a theory by being more precise in what we predicate of the entities mentioned in a theory.

[2] I have edited statement *p* for clarity.

These aspects, however, do not exhaust what it is to develop a genuine test. In order for a theory to be the subject of a severe and thus genuine test, it must do more than make a different prediction to those of its competitors. The development of different ‘empirical models’ designed to make predictions could be ‘tested’ to this extent. These models or ‘computation rules’ would no doubt become quite complex when trying to increase their scope and precision. However, this is not what Popper had in mind. Rather, a severe test of a theory involves a prediction that is unexpected and distinct from what is already entailed by the current state of background theory. In this sense, for a test to be severe, the prediction of the theory has to be ‘risky’. The unexpected prediction must be derived from an interesting new theory that is more general and/or more precise than what is entailed by the background theory. Importantly, if the prediction is incorrect then the theory can be considered falsified. If the interesting new theory survives a genuine test the theory can be considered ‘corroborated’. Popper believed that by the process of increasing falsifiability of competing theories in a scientific field, corroboration would be enough for theories to evolve in a way that approached the truth. Given two theories that range over the same subject area but make different predictions, the theory that survives a test is closer to the truth. This measure of a theory’s increasing approximation to the truth Popper termed ‘verisimilitude’.

Popper was interested in providing a formal account of his falsificationist programme, especially in terms of how his versions of logical probability and theoretical content might feature in a definition of verisimilitude. This project turned out to be quite problematic and the subject of serious criticism. Popper accepted much of the criticism, but argued that the formal project was meant to clarify his qualitative account, and thus was somewhat immune to the formal problems [Popper 1979: 371]. Although the qualitative approach is not beyond criticism [Maxwell 1972, Newton-Smith 1981: 52-67], we can still assess falsifiability as a way of demarcating what it *means* to be scientific. To this extent Popper provided some further advice:

i. It is easy to obtain confirmations or verifications, for nearly every theory – if we look for confirmations.

ii. Confirmations should count only if they are the result of *risky predictions*: that is to say, if, unenlightened by the theory in question, we should have expected an event which was incompatible with the theory …

iii. Every good scientific theory is a prohibition: it forbids certain things to happen. The more a theory forbids, the better it is.

iv. A theory that is not refutable by any conceivable event is non-scientific. Irrefutability of a theory is not a virtue of a theory (as people often think) but a vice.

v. Every genuine test of a theory is an attempt to falsify it, or to refute it. Testability is falsifiability; but there are degrees of testability: some theories are more testable, more exposed to refutation, than others, they take as it were greater risks.

vi. Confirming evidence should not count *except when it is the result of a genuine test of the theory …*

vii. Some genuinely testable theories, when found to be false, are still upheld by their admirers – for example by introducing *ad hoc* some auxiliary assumptions, or by re-interpreting the theory *ad hoc* in such a way that it escapes refutation. Such a procedure is always possible, but it rescues the theory from refutation only at the price of destroying, or at least lowering, its scientific status. [Popper 1963: 36-37]

We might note that, although Popper uses the term 'confirmation' in the above, it is clear that the sense of confirmation he is developing is not about gathering evidence to justify a theory. Although I agree with much of what these guidelines recommend, I will argue that points *ii, v* and *vii* are problematic. In fact, contrary to these guidelines, I claim that there is another kind of test that is not an attempt at falsifying a theory; that the test may recommend an *ad hoc* approach to the theory; and that we should consider the theory confirmed to some extent, even though it may not make a risky prediction. Most importantly, I will argue that a falsificationist needs something like this kind of test for a complete account of science.

2.3 Demarcating Pure from Applied Science

Popper thought that not only could the criterion of falsifiability demarcate science from non-science, it could also demarcate the pure from the applied sciences. In fact he thought that theories used within the limits of their applicability in the applied sciences were unfalsifiable. This seems to be because theories, in applied science, are treated merely as instruments of prediction: 'Thus a mere instrument for prediction cannot be falsified. What may appear to us at first as its falsification turns out to be no more than a rider cautioning us about its limited applicability' [Popper 1963:113]. Further on he says that '[i]n contrast to the highly critical attitude requisite in the pure sciences, the attitude of instrumentalism (like that of applied science) is one of complacency at the success of applications' [Popper 1963: 114]. The difference between the applied sciences and pure sciences is also to do with what the theories represent. The representational value of a theory for an applied scientist is not important if it is a good instrument of prediction, while on the other hand, for a pure scientist, one must think a theory an accurate representation of reality in order for it to be falsifiable. Although matters are not as clear cut as what is suggested here, there is a sense in which I agree with this. However, what I think requires further analysis is the nature of this unfalsifiability, especially in light of the 'success', of which the applied sciences are said to be so complacent. Unfalsifiability, I will argue, is rather beside the point if we are provided information from a test of an instrument.

Even if theories used in the applied sciences are considered unfalsifiable, we ought to demarcate those theories legitimately used from those theories that are used in the pseudo-sciences. Surely we want to demarcate the theory used in engineering design from the theory used in Freudian psychology?[3] Accordingly, one might reasonably ask for criteria for this second demarcation line. The issue here is not trivial, especially when we consider the scope of what might count as an applied science. Contenders may include the various forms of engineering: aeronautical, biomedical, chemical, civil, communication, electrical, electronic, fluid, geotechnical, information, instrument, marine, mechanical, mining, petroleum, software and structural to choose just a few. Nor need we limit the scope of applied science to engineering; the fields of agriculture, geology, biology, medicine, chemistry as well as the physical and environmental sciences have many applied and pure forms. It is true that there are many pragmatic considerations unique to each field. However, what I take to be the important aspect, common to all, is the role of a design or model based on, at least in part, established theory, initial assumptions, and test work. To this end I argue that design theory, although it can be treated as unfalsifiable in the way that Popper suggests, it is also testable (in a sense to be developed).[4]

According to Popper, the practitioners of the applied sciences don't in general set out to falsify a theory.[5] In so far as the applied sciences view testing, it is a method of optimising applicability. If it is discovered that a theory is not applicable in scope or precision, an applied scientist views the test as establishing just that. Lack of applicability might then send the applied scientist back to the 'tool box' to find a better instrument, or indeed modify a theory so that it is suitable for the task. On the other hand, a pure scientist sees the failure of the theory to apply in a context as a falsification of the theory. This then provides an opportunity to go back to the theoretical realm and reconceptualise the problem. Not only should the pure scientist try to solve the known problems, but also, the newly developed theory should make unexpected predictions that could feature in a severe or genuine test. For Popper, the distinction between the applied and pure sciences can be understood as a difference in attitude that has methodological implications. Applied scientists (with instrumentalists) see theories as tools, while pure scientists see theories as representations of reality that can be true or

[3] Whether or not Freudian theory is testable has been the subject of some debate [Grünbaum 1984]. I use it because Popper believed that it was not falsifiable in principle. The current intelligent design theory might be a modern example, or the enormous range of 'new-age' theories that deny the relevance of evidence garnered from randomised control group testing or double blind clinical trials.

[4] The idea of a testable but not falsifiable theory is not new. Donald Gillies has developed such an idea and how it might be related to a theory of confirmation [Gillies 1993: §10.5]. My theory is different to the one offered by Gillies.

[5] Popper acknowledges that applied scientists do test for errors [Popper 1979: 353n]. Though, as we will see, what this testing amounts to for a falsificationist is problematic.

false.[6] According to Popper, if we just followed a strategy of searching for applicability, pure science would stagnate.[7]

We will soon see there are difficulties with this rather neat characterisation. One obvious problem is that it only refers to the success of the pure sciences. It seems that, although pure science would progress if it followed the falsificationist model, it would stagnate if it followed the methodology of the applied sciences (or instrumentalism). But the applied sciences have also been successful, especially in the invention and optimising of experimental technology (and artefacts more generally). In fact, what is particularly impressive about the applied sciences is the resourcefulness and creativity displayed in the design of experiments used to test the theories developed by the pure scientists. It is not at all clear that a falsificationist can account for this success.

To appreciate how serious this problem is, let us consider in more detail how Popper characterises the instrumentalist. Although Popper thinks the instrumentalist, with the applied scientist, can perform a 'test' of sorts, it is different to the kind of genuine test the falsificationist is advocating. 'The way in which computation rules are *tried out* is different to the way in which theories are *tested*' [Popper 1965: 111]. It is not clear what *trying out* a theory amounts to, although it is clear it is a different kind of activity to a *severe test* of the theory:

> Nothing sufficiently similar to such tests exists in the case of instruments or rules of computation. An instrument may break down, to be sure, or it may become outmoded. But it hardly makes sense to say that we submit the instrument to the severest tests we can design in order to reject it if it does not stand up to them: every air frame, for example, can be 'tested to destruction', but this severe test is undertaken not in order to reject every frame when it is destroyed but to obtain information about the frame (i.e. to test a theory about it), so that it may be used *within the limits of its applicability* (or safety). [Popper 1963: 112-3]

It is a little obscure what the precise point is that Popper is making here. Airframes are tested, and if one of them fails what is considered the worst-case flying conditions, then (we hope) every air frame failing such a test will be rejected. So in some sense instruments can be rejected because of a test. Given that Popper is contrasting 'trying out' an instrument with the severe test of a theory, it seems reasonable to ask what this trying out amounts to if it is not a severe test. The issue is not clarified by Popper's own

[6] For Popper these representations have to be described in sentences that can be true or false.
[7] There are other accounts that contest this claim. Thomas Kuhn [1970] has suggested that science is a form of non-truth-directed puzzle solving. Nicholas Maxwell argues that even if we accepted that the aim of science as an attempt at getting close to the truth, Popper's methodological rules by themselves do not help achieve this aim [Maxwell 1972].

interpretation of what it is to test an instrument. Especially problematic is the use of the term ‘severe test’ that determines the limits of an instrument’s applicability. Maybe this term should have been put in scare quotes? If a severe test determines the applicability of the instrument (the air frame) then, it seems, there is no methodological difference between pure and applied science.

The issue is clarified somewhat a little further on.

> We may sometimes be disappointed to find the range of applicability of an instrument is smaller than we expected at first; but this does not make us discard the instrument *qua* instrument – whether it is a theory or anything else. On the other hand a disappointment of this kind means that we have obtained new information through refuting a *theory* – that theory that implied that the instrument was applicable over a wider range. [Popper 1963: 113]

Although this doesn’t help clarify the distinction between trying out and severe testing, it now seems clear that we get two things from a severe test. We get the information about the applicability of a theory, construed as an instrument (or anything that might be construed as an instrument), as well as the information that refuted the (different) theory that is construed as something that can be true or false.

Notwithstanding the points made above, the central ideas Popper is driving at are quite plausible. Firstly, we don’t try to falsify an instrument; this seems to be a category mistake. The terms ‘true’ and ‘false’ apply to sentences or propositions not tools or instruments. Secondly, there are two kinds of information derivable from a test: the information to do with applicability, and the information that might refute a theory. Given these plausible observations, we might still note that it is no contradiction to say that we can test an instrument. In fact it seems we often do test instruments, and not just the theories about them. When I test a knife to see how sharp it is, a falsificationist will interpret this as an attempt to falsify one of at least two competing theories regarding how the knife will behave in the test. But we are also testing the instrument – the knife – to see how sharp it is. The information garnered from the testing of the knife is important if we wish to use the knife effectively. This, I think, is the key difference between a severe test of a theory and what might constitute the trying out of an instrument. Trying out an instrument is ‘testing’ the instrument’s applicability. Trying out the instrument provides us with information on how to use it.

Rather than use the term ‘trying out’, I think it more helpful to call this kind of testing an ‘informative test’ (because of the information it provides). The severe test may provide information that can falsify a theory, but the informative test is more useful. It provides information on how to use the instrument. Popper could try to explain the informative test as a form of severe testing, but if he were to do this, it comes at the price of destroying the methodological distinction between pure and applied science. Or so I will argue. But first, let us turn to an example of an informative testing used in the applied sciences.

Let us suppose that a new theory is proposed that requires testing. Let us also suppose that the testing of this theory requires reasonably accurate predictions of the pressure profiles of a fluid flowing at various velocities through a series of pipes.[8] Now an applied scientist needs some design theory in order to build this instrument to the specifications required. To do this, the applied scientist might identify the relevant background theory. A reasonable starting point here would be to note a 'classical' understanding of force as a function of pressure difference ($p_1 - p_2$) and area (A).

1. $F = (p_1 - p_2)A$

We also have available the successful theory that suggests that frictional force developed by the turbulent flow of a fluid is a function of the surface area (PL) in contact with the fluid, the velocity of the fluid (v) squared, and friction resistance at unit velocity per unit area (q).

2. $F_f = qv^2PL$

Given that pressure difference can also be measured in terms of vertical length of fluid (h_f) of a certain density (w) we obtain:

3. $h_f = (p_1 - p_2)/w$

It is then straightforward mathematics to derive the design principle for pressure loss (h_f) developed in a pipe through which a fluid is flowing.

4. $h_f = qPLv^2/wA$

Although this formulation – known as the Darcy formula – is quite suitable, the Chezy formula (derivable from it), is sometimes more convenient.[9]

5. $v = C\, m^{0.5}\, i^{0.5}$

In the above formula, the Chezy coefficient (C) represents frictional properties of the fluid and pipe work. The hydraulic radius (m) is a function of the perimeter and area of the pipe, while pressure loss (i) is now relative to a unit of pipe length. It now seems

---

[8] The following derivation of the design principles is informed by Douglas [1986: §10.5].

[9] For the derivation assume that $i=h_f/L$, $m=AP$, $C=(w/q)^{0.5}$.

possible for our applied scientist to design the pipe work so that the required pressure profiles can be reliably predicted.[10]

Although the theories used to develop the Chezy formula have been well tested and successful, testing instruments designed on the basis of the Chezy formula reveal that the accuracy of the predictions is compromised relative to the velocity of the fluid. From the perspective of the applied scientists, the accuracy of the design may still be within the tolerable limits for the experiment to work. On the other hand, from the perspective of a falsificationist, the difference between the prediction and test results may be considered to falsify the Chezy formulae and one or more of the theories that underpin it. Still, according to a falsificationist, it is rational to use a severely tested, though strictly speaking false theory, if that is the best-tested theory we have.

Interestingly, an applied scientist need not be so complacent about this result. As mentioned, testing an instrument need not only falsify a theory, it can provide information about the applicability of the theory. In fact there are often opportunities to develop more applicable theories based on the information obtained from testing an instrument. Thus we get 'modifications' to the Chezy formula. For example the Hazen-Williams formula:

6. $v = 0.82C_1\, m^{0.6}\, i^{0.54}$

This formulation is a better fit to the information garnered from testing the instruments. Notice the new coefficients and adjustments to the theory are determined by testing the instrument, and, as a result, this formula is not severely tested in Popper's sense. In fact it makes no radical new prediction outside the context of its discovery. It is not a result of a re-conceptualisation of the problem and attempted refutation. Indeed it looks like an *ad hoc* adjustment of a theory to fit the data – the precise methodology the falsificationist abhors. So now we have two theories – one severely tested (Chezy) and an *ad hoc* theory that is not genuinely tested (Hazen-Williams). According to Popper's guidelines (*ii* above) 'confirmations should count only if they are the result of *risky predictions* …' [Popper 1963: 36]. So it looks as though the Hazen-Williams formula is unconfirmed. But here comes the rub. What if our experiment requires a more accurate design than the Chezy formula can provide? How might we be justified, from a falsificationist's perspective, in using the untested Hazen-Williams formula?

One response might be to say that, in so far as it does make some predictions we can test the Hazen-Williams formula. This test might not meet the standards of a genuine or severe test, but the design theory is, to this limited extent, testable. This response just turns a dilemma into a difficult decision. We now have to decide between a design

[10] There will be much more to it than the use of just this design theory, but this should be enough of a sample to get the general idea.

theory based on successful, severely tested and falsified theory, and design theory that, if based on anything, is weakly tested and thus of very limited success (and that because of its *ad hoc* development provides no improvement in verisimilitude). What we require is some reason to choose the Hazen-Williams formula, and given that a falsificationist cannot admit that there is any justification for believing that it *is* an accurate instrument of prediction, they risk choosing the wrong theory.[11] The problem disappears if we admit the relevance of an informative test. The justification for the use of the Hazen-Williams formula is based on information derived from testing the relevant instruments, and by comparison, is a more accurate predictor than the Chezy formula.

I am not suggesting that it is just this type of testing and *ad hoc* adjustment of theory that is indicative of the applied sciences. Rather, the point of the example is to highlight the informative nature of a test that might, in fact, be a severe test. From the perspective of an applied scientist, a severe test can provide information about the rate of false positives and or false negatives, or provide a good determination of the margin of error to be expected in a theory's application. This is good information to have when applying the theory in the future. And as we have seen, this information might even suggest a better predictive tool to help in the design of a more precise instrument. From the perspective of the falsificationist, this information only serves to corroborate or falsify a theory.

It is clear that Popper does think that severe testing can provide information about an instrument's applicability. If this means more than just determining where an instrument is not applicable, a falsificationist needs to claim that surviving a test (corroboration) provides 'information' about where the instrument can be applied successfully. To avoid accusations of inductivism, the falsificationist will need to add the qualification that the corroboration of the theory itself does not justify any predictions implied by the information. Unfortunately this qualification creates a problem. To see why, consider what would happen if similar instruments failed to perform according to the information provided by the corroborating severe tests. The appropriate response is to claim that we have been *misinformed*. For whatever reason, the original tests did not provide the relevant information about the predicted behaviour. Information, in this context, carries with it its own justification, otherwise it is not information. If I am supplied information on how an instrument will behave, I am also supplied some justification for believing that the instrument will behave that way.[12] But if corroboration has nothing to do with the justification of a prediction, it now seems mysterious where the justification required of the information comes from.

[11] This point is different to one developed by Wesley Salmon [1988]. Salmon argues that corroboration can't provide a rational basis for pragmatic prediction. My point is that the falsificationist methodology might positively misguide us when it comes to deciding on a theory for pragmatic prediction.
[12] This concept of information owes much to Fred Dretske's work on the subject, especially Dretske [1981].

It could be argued that I have misrepresented what information is for a falsificationist. Maybe, for a falsificationist, information has no justified predictive content. The predictive content is entirely conjectural. However, even if all that is true, we now have no reason to believe that this kind of 'information' will have any relevance to how other instruments similarly built will behave in the future. Or, indeed, why the information regarding an instrument's accuracy, measured prior to its use, has any relevance when used.

This problem is evident in Popper's own example of what it is to use a theory as an instrument:

> For instrumental purposes of practical application a theory may be used *even after its refutation*, within the limits of its applicability: an astronomer who believes that Newton's theory has turned out to be false will not hesitate to apply its formalism within the limits of its applicability. [Popper 1963:113]

Is Popper advocating that for practical purposes a falsificationist can provide a way of *justifying* where and when to apply the 'formalism' of the theory? Although the text seems to indicate we can, the answer has to be that we can't. If applicability is just an assessment regarding the past application of the theory, then any success cannot justify the future use of the theory within these historical limits. What then is the theory underpinning Popper's prediction regarding the astronomers? Some might suggest, with Hume, that it is just habit that explains the lack of hesitation on the part of the astronomers, but presumably Popper would deny this [Popper 1963: 42ff.]. Nor is the astronomers' lack of hesitation explained by the fact that they are trying to falsify a theory. Apparently they appreciate that the theory being applied is already falsified. I suppose Popper could say that the predicted cavalier approach displayed by our astronomers might be because they *accept* that past testing of the theory warrants its future use within the limits established by such testing. This may explain the behaviour of our applied scientists, but the accepted methodology can't feature as the likely explanation for the success of this form of science. The falsificationist thinks that there is no evidence to support such a methodology and thus it is very unlikely to be the correct explanation for the successful use of the formalism of Newton's theory.

One tempting response is to try and bypass these issues by supplementing the genuine test that demarcates science from non-science with the following characterisation: pure science is about developing theories; applied science is about applying them. This might allow us to explain a theory's success in terms of severe testing, while allowing us to demarcate pure from applied science. On this construal, the issue of justification might seem to fall out of the picture: applying theory is just what applied scientists do. In order to do this, applied scientists are guided by rules, conventions, techniques and the like. Unfortunately, this only delays the issue. The rules, conventions or techniques, we might presume, are not arbitrarily chosen. Thus we might hope to consider their justification. Yet this characterisation is also too neat to be true. Pure scientists will need to develop theories cognisant of the results of previous tests and with an eye to how one might test a newly developed theory. As such, a pure scientist when developing a theory must take into account its applicability.[13] On the other side of the ledger, applied scientists do, to some degree, develop something like theories or models that they test. Moreover, given the enormous number of ways a theory might be applied to get a nil result, the development of a successful test (whether falsification or corroboration) is a significant achievement in itself. Despite the use of 'unproblematic' background theory in the design of a new experiment (or artefact), a degree of ingenuity in the design and an enormous amount of error checking is usually required. The design of a test is much more than the application of theory.

Another approach, suggested by Mario Bunge [1966], characterises a technological forecast as always in the interests of some predetermined outcome. Here the applied scientist will adjust the assumptions and the instrument itself to achieve that outcome. On the other hand Bunge claims a pure scientist is only interested in understanding the world, and will let the experimental dice fall as they may [Bunge 1966: esp. 341ff.]. However, Bunge has failed to notice that it is the applied scientists who have to design the experiments to test the theories that pure scientists provide. If this is right, it implies that the applied scientists, contrary to his claims, are the true keepers of impartiality.

One response might be to argue that there are in fact three kinds of scientist, pure, experimental and applied. Bunge could then be read as contrasting the pure and applied sciences, leaving the experimentalist to do the impartial testing. Bunge does not make this distinction, and although I think it could be made, more analysis would be required in order to understand the differences and relationship between the three kinds of science. More importantly, the introduction of a third kind of science to demarcate would hardly help a falsificationist.[14]

Having provided some detail to the complex relationship between applied and pure science let us review what is at stake. The genuine testability of a theory is meant to explain the success of science, as well as demarcate pure from applied science and pseudo-sciences more generally. Applied science, in so far as it does not severely test, is a non-science. Pure science, in so far as it does severely test its theories, explains why it

---

[13] This may mean that the 'results' need to be reinterpreted in light of the new theory, or even questioned. However, previous results cannot be ignored in the development of a theory.

[14] There are other attempts at demarcating applied and pure science: Joseph Agassi [1980], Paul Quay [1973] and David Resnik [2000] all identify additional demarcation principles, but neglect to highlight how the pure sciences depend for their results on the applied sciences. An exception is Daniel Rothbart [1998] who reviews how theories might be applied in the lab using Popper's epistemology (and metaphysics). Much of Rothbart's analysis I agree with, but only because it seems to provide a *justification* for the information we receive from laboratory instruments.

is successful. However, both pure and applied sciences have been successful. So we are owed an explanation for the success of the applied sciences. It hardly seems an explanation that the applied sciences are successful because they do not genuinely test. Popper also has to make good on the claim that the different methodology of the applied sciences, if adopted by the pure sciences, would lead to stagnation.

It is clear that Popper thinks that what would cause the pure sciences to stagnate has something to do with how the applied sciences treat certain theories. Applied sciences, it seems, treat theories as unfalsifiable when considered as instruments, and it is this attitude that would stagnate the pure science. But regardless of the effect that this approach might have on the pure sciences, we might ask, what is the purpose of unfalsifiability in the applied sciences? If this distinctive approach explains the success of the applied sciences, then Popper's demarcation principle is inadequate. We now have some success in science that does not depend on a genuine test. If on the other hand treating theories as unfalsifiable also limits the success of the applied sciences, it is not clear why we should consider unfalsifiability methodologically relevant to the applied science. Reversing the perspective, we might easily identify some bad or irrelevant practices in the pure sciences and then claim, if the applied science were to follow the methods of the pure sciences, the applied sciences would stagnate. Presumably, those in the pure sciences could just deny that those bad or irrelevant practices were any part of the methodology of the pure sciences. But if the same response is granted to the applied science, then they too can reject Popper's interpretation of their methodology.

There may be another approach. Popper might be able to claim that the applied sciences can 'severely test' theories for applicability (rather than severely test in order to falsify a theory). Notwithstanding the difficulties a falsificationist faces in providing a viable account of applicability, it is not clear why this approach would result in any methodological difference between the two kinds of science. All testable theories have to be applicable in principle, and it is applicability or lack of it that decides if a theory is corroborated or falsified. Granted, this method *might* explain the success of the applied sciences. But then how would this method limit the progress of pure science? The dilemma for falsificationism seems to be this: if applied science is a non-science, we get to explain why the pure sciences would stagnate using the applied approach. But then we can't explain the unique success of the applied sciences. On the other hand, if the success of the applied sciences can be explained with just the resources of the severe test, we can't distinguish, methodologically, between applied and pure science, and thus why the pure sciences would stagnate using the applied methodology.

If the demarcation principle supplied by the falsificationist is not an adequate explanation of the success of the sciences then we are owed something more. We need an account of what it is that is methodologically unique about the applied sciences and why that would impede the development of the pure sciences. Furthermore, we need an account of the role of this methodology in the complete explanation of the applied sciences, and how this might relate to its success. While I think the informative test helps falsificationism meet these requirements, it seems falsificationism can't make sense of this kind of test.

2.4 *Some ad hoc Problems for Falsificationism*

So far, we have looked at Popper's criteria for demarcating science from non-science, and then pure science from applied science. The principal criterion – falsifiability – Popper equates with testability. I have argued that there is a second type of test (or aspect of a severe test) that might be used to *justify* a theory's applicability. Without this kind of test we cannot explain the total success of the sciences. In this section I will argue that there is a tension generated in Popper's solution to the Duhem-Quine problem.[15] This tension is avoided if the falsificationist allows for the informative test.

The concept of *falsification*, as explained by Popper, seems rather naïve when we consider the cut and thrust of actual science.[16] Imre Lakatos [1978], for instance, has noted that when we look at actual scientific practice, there is usually a central 'hard core' of hypotheses not open to falsification. Protecting this hard core is a belt of potential excuses. If the hard core is threatened, then we can use the rules of excuse-making (the 'negative heuristic' as he terms it) to explain away the falsifying evidence. One reason why this might be considered an acceptable practice is that theories often require time to develop. If you falsify a theory before its time you may inhibit scientific development. Hence, we should not so much assess each individual theory, but a series of theories as they develop. Lakatos terms this development a 'research programme'. Making fruitful predictions and explaining past failures in terms of the developing theories count as a growing research programme. If, on the other hand, the negative heuristic is appealed to more often than not, the research programme is degenerating. As mentioned, it is legitimate to use the negative heuristic to protect the central aspects of the research programme, but only if there is a prospect that the diversion might, at some later stage, be turned into a success.

Paul Feyerabend [1978: Ch. 8] contrasts the difference between Popper and Lakatos in terms of their attitude to the deployment of *ad hoc* hypotheses.

[15] There are in fact two theses, the Duhem thesis and the Quine thesis, the latter being more global than the former. The sort of problems identified can be generated by either thesis. For more analysis here see [Gillies 1993: esp. ch. 5]

[16] There is another sense in which falsification is too simple. We might notice that there are many falsifiable statements that we would not claim to be science. For example, 'whenever I run out of milk, I always immediately go to a shop to get some more' is a falsifiable yet hardly a scientific generalisation. Larry Laudan [1983] has used this type of criticism to some affect.

> Popper: New theories have, and must have, excess content which is, but should not be, gradually infected by *ad hoc* adaptations.
>
> Lakatos: New theories are, and cannot be anything but, *ad hoc*. Excess content is, and should be, created in a piecemeal fashion, by gradually extending them to new facts and domains. [Feyerabend 1978: 93-94]

Even if Feyerabend is overstating the case, there is evidence for this thesis. For instance, Popper argued that the 'complementarity principle' in quantum theory was an *ad hoc* interpretation that protected the theory from falsification (by allowing what seemed to be an inconsistency). As a result;

> the principle of complementarity has … remained completely sterile within physics. In twenty seven-years it has produced nothing except some philosophical discussions, and some arguments for the confounding of critics (especially Einstein). [Popper 1963: 101][17]

From Lakatos' perspective, this kind of protection of a theory could be justified, but only if it offered the prospect of further success: 'It may be rational to put the inconsistency into some temporary *ad hoc* quarantine, and carry on with the positive heuristic of the programme' [Lakatos 1970: 143]. The example Lakatos uses is Bohr's 'correspondence principle' which functioned to protect the programme against an inconsistency between quantum and classical physics, while also offering novel hypotheses that could be tested. Of course it is notoriously difficult to assess the *prospective* success of a theory (see note 17), so what counts as a rational *ad hoc* move may have to be determined in hind-sight.

The nature of the *ad hoc* problem is a deep one for falsificationism, and is at the core of the Duhem-Quine thesis.[18] According to this thesis, in any falsifying test situation, there is an enormous range of auxiliary theories that are required to interpret the results, not least of which is the theory behind the functioning of the experimental technology. In any falsifying situation the scientist does not know which of these theories ought to be falsified. This provides considerable scope for saving a theory.

A telling example is the discovery of the particle nature of the electron by J.J. Thomson. Prior to Thomson's important experimental investigations, there were falsifying experiments for both the wave and particle theories accounting for cathode rays. There was, for instance, evidence that cathode rays were able to 'carry' a charge, a point against the wave theory. Yet it seemed that cathode rays could penetrate through solid

[17] Interestingly this programme turned out not to be as sterile as Popper originally thought [Popper 1963: 114]

[18] There is an excellent discussion of this thesis, and of its correct formulation, in Gillies [1993].

objects without puncturing them: a point against the particle theory. Moreover, it was discovered by Hertz that cathode rays could not be deflected by an electrical field: another falsifier for the particle theory. Thomson's solution was to introduce, in an *ad hoc* way, one testable hypothesis and one untestable hypothesis in favour of the particle theory. Firstly, he hypothesised that electrons are small enough to penetrate through thin material objects (*ad hoc* and untestable at the time); and secondly, he blamed Hertz' experimental design rather than the particle theory (*ad hoc* but potentially testable). It seems that the effect of the electric field in Hertz' experiment was reduced by the presence of ionised gas in the cathode tube. By improving the vacuum in the cathode tube – removing most of the ionising gas (no small task) – the electrical field was able to divert the cathode ray, turning a falsifying experiment into a marvellous success. This improved technology allowed Thomson to make important experimental discoveries on the nature of an electron; discoveries which won him the 1906 Nobel Prize for physics. The irony is that his son (George) went on to win a Nobel Prize for discovering the wave-like nature of the electron![19]

Thomson spent a lot of time and effort getting a sufficient vacuum to test his theory. We might ask ourselves, had he failed in the first instance should he have given up? An even more pertinent example is Michael Faraday. Possibly one of science's most successful experimental scientists, Faraday was notoriously dogmatic. More often than not, his experiments would fail before he finally modified the set up so as to discover the phenomena he was after [Tweney 1985]. Despite what looks like an *ad hoc* approach, his theory of electromagnetism and its relation to electricity, matter and light resulted in many significant discoveries.

Popper's response to the Duhem-Quine problem is to claim that what is rejected in any falsifying test is 'the theory *together* with that background knowledge … parts of which, if other crucial experiments can be designed, may indeed one day be rejected as responsible for the failure' [Popper 1963: 112]. But this seems to just succumb to the Duhem-Quine problem. Popper can now be interpreted as advocating that it is scientific to consider designing a genuine test of any hypothesis responsible for a theory's falsification. But now it is difficult to know when to give up on a theory. Certainly we can't decide with the help of an inductive argument based on past failures.

Key, then, is some analysis of what it is to develop a theory in an *ad hoc* way. Popper, throughout his writing, outlines several kinds of *ad hoc* move.[20] For instance, we may try and save a theory by introducing, in an *ad hoc* way, an untestable hypothesis that reinterprets the results of a test [Popper 1974: 986]. He also notes that an *ad hoc*

[19] For a more detailed account of the controversy with Thomson's resolution see [Davis & Falconer 1997: esp. 122ff]

[20] Popper also labels the *ad hoc* move as a 'conventionalist twist' or 'conventionalist stratagem' [Popper 1963: 37]. Other rescuing operations might include denying the results are 'scientific' or calling into question the objectivity of the experimenter [Popper 1965: 81].

hypothesis may become testable, and thus at this later stage treated as an auxiliary hypothesis. An auxiliary hypothesis can be introduced for *ad hoc* reasons but they just so happen to be testable [Popper 1974: 987]. Presumably auxiliary hypotheses thought to be testable can also turn out not to be testable (not all tests of hypotheses give definitive results) and thus can be found to be an *ad hoc* hypothesis after all. We might also, in an *ad hoc* way, make certain corrections to our measuring instruments, changing the results to suit a theory [Popper 1965: 80-1]. Alternatively, in order to escape certain contradictions, we can take, as an *ad hoc* manoeuvre, an instrumentalist attitude to a theory [Popper 1963: 101]. Interestingly, the *ad hoc* move need not exclude testability. We might introduce a series of *ad hoc* yet testable theories that do not pass any experimental test [Popper 1963: 244]. We could, for example, introduce in an *ad hoc* way, an auxiliary hypothesis to save a theory from falsification. If the auxiliary were falsified, we could introduce, *ad hoc*, another auxiliary that explains the failure. If in turn it were to fail its test we could introduce, again *ad hoc*, another theory that explains the problem and that, once again, is ready for falsification. And so it goes. In fact Popper asserts that 'in this case, we should feel that we were producing a sequence of theories which, in spite of their increasing testability, were *ad hoc*, and that we were not getting any nearer the truth' [Popper 1963: 244].[21]

Testable or not, hypotheses introduced or adjusted in an *ad hoc* way are attempts to protect a theory from some of the evidence so far gathered. Popper later modified his demarcation criterion to accommodate some amount of this kind of *ad hoc* manoeuvring. Apparently we need to protect our theories in order for them to develop [Popper 1974: 984]. So there seems to be legitimate and spurious uses of the *ad hoc* move. Although some degree of legitimate protection is allowed, we can't keep introducing explanations for anomalies indefinitely. In order to protect against this spurious use of the *ad hoc* move, Popper suggests that in a theory's development it must have passed some genuine test – enjoyed some success. Although vague, this sounds like good advice. Unfortunately the advice creates a problem for another part of his theory.

Popper sometimes refers to hypotheses that are *easily* testable, suggesting that what makes a hypothesis easy to test is a matter of training in a technique [Popper 1965:99]. We will consider what this means in more detail in the next section. For now, let us consider an example of an easy test. Although we can never be sure of the statement 'this glass contains water' it is, we are assured, easy to test. Presumably this is because of the dispositional properties involved [Popper 1963: 387-8]. Maybe our understanding of water is such that we believe all samples of water have certain properties without

[21] It is not clear to me why we should, by falsificationist standards, feel that we are *not* approaching the truth after a series of failures. Maybe we are looking for a very small or rare effect; say for example, a gravitational wave. It seems reasonable to expect that there could be a long history of failures before the effect will be discovered. Of course we could get lucky very early on in the search, but there is some reason to feel that a long series of failures might also be getting us closer to the truth.

exception. If the properties of water were less determinate then the test would not be easy. We would have to deal with a range of false negatives or false positives or both. In order to perform this 'easy test' then, we must have tested those theories describing the properties of water. When testing the properties of water we need to propose 'basic statements' that contradict the predictions of the theory that describes the nature of these properties [Popper 1965: 33, §28]. As all statements are theory-laden, so are the basic statements required to test the predictions of a theory.[22] If these basic statements are continually falsified, then presumably the theory that accounts for them has not enjoyed any success. But testable theories that have had no success are spuriously *ad hoc* or will involve spurious *ad hoc* moves.

This point may look trivial, but Popper has defended the notion of falsifiability on the possibility of just this kind of easy test. Apparently Newton's theory is falsifiable at least to the extent that the moon might fly off at some tangent, or apples might rise up from the ground and start dancing around their branches [Popper 1974: 1004ff.]. Presumably the theories that account for these radical predictions (the basic statements required to test Newton's theory) have not enjoyed any success. But it now seems that the theories (implicitly) used in the defence of the falsifiability of Newton's theory are spuriously *ad hoc*. For, at the least, we need a series of spurious *ad hoc* moves to explain all the previous and ongoing refutations of these alternatives to Newton's theory, just so that we can plausibly suggest such radical predictions. To test Newton's theory in this easily testable sense we require an unscientific theory by falsificationist standards. Spurious *ad hoc* moves are thus required to defend falsifiability. Maybe this result is not that devastating for a falsificationist, but it does indicate an uneasy alliance between the need for spurious *ad hoc* moves and an easily testable theory.

It should be clear that the informative test avoids this problem. Some tests (or aspects of tests) are not attempts at falsification, but are designed to provide information. Although we can, there is no need to propose a counter theory when testing an instrument to obtain this kind of information. As Popper noted, falsification is not an appropriate term when accounting for this kind of test (trying out) of an instrument. Newton's theory can be easily tested, at least to the extent that experiments designed on the basis of this theory can be shown to be very accurate in what they predict in well understood applications (the informative test). But if this is the solution, falsificationism relies on an understanding of a test that has little to do with a severe test. Maybe Popper can explain an easy test as a form of severe testing. But this explanation, I suspect, would make the severe test too permissive, and thus would expand what we hope to count as a pure science. If, on other hand, performing an easy test is just a matter of training in a

[22] It doesn't really matter that no one holds the theory that accounts for the required basic statements, only that it is believed that such a theory is possible. If it is not possible to formulate a theory that might account for the basic statements, what is being proposed in the test is theoretically impossible. This certainly makes the test easy, but for the wrong reasons.

technique, and this technique is not to do with developing a severe test, it looks like the easy test is not scientific.

2.5 Corroboration and Reproducibility

It might seem that what it is to easily test a theory is related to the concept of reproducibility. An easily testable theory might be so because it is easy to reproduce. If reproducibility can be accommodated within the falsificationist programme, then maybe we have enough to account for what makes a test easy (or hard) without appealing to the informative test. I will argue, however, that establishing the reproducibility of a test relies on an informative test, or once again we lose the distinction between applied and pure science. Both outcomes are unacceptable for a falsificationist who believes they can provide an accurate account of good scientific practice.

Usually a single falsifying observation is not enough to reject a whole theoretical system.[23] In order to ensure that the test was not an accident, evidence has to be reproducible. On the face of it, if reproducibility is important to falsification, then we need some reason to believe that the test is reproducible. How might we account for this belief? It is hard to see how corroboration might help here; Popper consistently claims that corroboration has nothing to do with justifying further predictive success of a theory. The problem is that reproducibility involves a modal claim regarding the success of an experiment in different contexts. We can express this modal aspect with the use of a subjunctive or counterfactual conditional: if certain experimental conditions (A) were to be the case then certain result (B) *would* be observed. Corroboration however, is a measure of past success and gives no justification for success of a test in other possible contexts. The problem, it seems, is that Popper explicitly uses a rich sense of reproducibility as a necessary condition for a test.

> Only by such repetitions can we convince ourselves that we are not dealing with an isolated 'coincidence', but with events which, on account of their regularity and *reproducibility*, are in principle inter-subjectively testable [Popper 1965: 45; my italics].

A little further on, Popper defines a *physical effect* as:

> [T]hat which can be regularly reproduced by anyone who carries out the appropriate experiment in the way prescribed [Popper 1965: 45].

Both these statements include modal claims about the reproducibility of the experiment necessary for an inter-subjective test. Given that we are not justified in believing claims

[23] The fact that scientists sometimes do claim that an experiment is repeatable after one or two tests is explored further in Cartwright [1999].

about the future (or the past), how is a scientist to know that a physical effect is reproducible rather than a coincidence? Repeating a test might help settle the matter if the results of further testing conflict with the original findings, allowing us to decide that the original test was an isolated coincidence. However, it seems unlikely that, even after repeated testing, a falsificationist could ever be ‘convinced’ that he is not actually dealing with a series of isolated coincidences, and that he has failed to establish that the test is inter-subjectively reproducible.

One solution might be to argue that certain rules outlining the evidential weight of test results (basic statements) can do the work of reproducibility here. Once we have established what it *means* to falsify a theory, there is no need to continually reaffirm this. Once we have decided upon the amount of testing required we have enough to call a test result reproducible. The decision may depend on a range of factors informed by what we accept as the relevant, though potentially problematic, background knowledge and assumptions. Thus we need only decide what constitutes sufficient evidence given the type of test involved, and what we are prepared to accept (without justification) is the case.[24]

Yet it seems Popper is sensitive to the potential arbitrariness of such decisions.[25] Popper acknowledges that there is some relationship between our experience of the world and a basic statement:

> [T]he decision to accept a basic statement, and to be satisfied with it, is causally connected with our experiences.

But Popper immediately takes back what looks like a promising lead.

> But we do not attempt to *justify* basic statements by these experiences. Experiences can *motivate a decision*, and hence an acceptance or a rejection of a statement, but a basic statement cannot be justified by them – no more than by thumping the table. [Popper 1965: 105]

This qualification seems to raise important questions. What is it for experience to ‘motivate a decision’? If motivation has nothing to do with justification, why should we care what particular experience motivates a decision? We may as well let any coincidence motivate the decision that we have discovered a reproducible effect. To

---

[24] Something like this solution to the problem of what constitutes a reproducible effect for a falsificationist is offered by Alan Musgrave [1975]. At some point we decide we have sufficient evidence that the type of test is reproducible. Unfortunately, Popper also claims that his theory can explain why there are diminishing returns for repeated testing. But if Popper is to buy into Musgrave’s solution he loses this explanation. This is not a serious problem, according to Musgrave, who argues that we should not hanker for a law of diminishing returns.

[25] See for instance Popper [1965: 204-5].

avoid this problem, Popper argues that reproducibility depends upon the precision of our *technique of measurement* [Popper 1965: 124, 202, 204]. Unfortunately, this leads straight back to our informative test. To see this more clearly, we will have to consider, in a little more detail, the problem of what counts as a reproducible effect, and how it features in Popper's theory.

Popper acknowledges that scientists could potentially justify any result they want just by virtue of the selective sampling of data. Probabilistic theories are particularly susceptible to this problem. Yet scientists frequently use probability; how is it that they can use probability theory while retaining falsifiability? According to Popper, the tentative answer is that:

> Frequent repetition leads to results with relative frequencies which, upon further repetition, approximate more and more to some fixed value which we may call the probability of the event in question. [Popper 1965: 198-9]

This result makes it possible to falsify any conjectured estimate of probability. Unfortunately, there are several problems with the above approach. The theorem that the probability will converge – Bernoulli's theorem – is not universally applicable. How do scientists know they are dealing with a convergent, that is reproducible, data set? Moreover, how do scientists determine what counts as a large enough sample size for a calculation of the probability; or, indeed, what that approximation is? Popper's solution is to show how these three problems reduce to one – a decision as to what is an acceptable probability of error. According to Popper, if ΔP is the acceptable range for a successful measurement, and ε the degree of error, or the probability of not ΔP, then 'ΔP can … be taken to correspond to the interval of precision … which depends on our technique of measurement' [Popper 1965: 202]. The decision to neglect ε occurs when changes in its magnitude do not significantly affect what we determine as ΔP. Importantly, this relationship is dependent upon the sample size we choose. By choosing the correct sample size it looks as though we have a way of solving the problem of when ε is small enough.

Apart from choosing a sample size we might wonder what it is for ΔP to depend upon a technique of measurement. Popper stresses that ΔP is not a clear cut range but is fixed around 'un-sharp bounds' or 'condensation bounds' as he terms them, even so:

> If we measure a magnitude many times, we obtain values which are distributed with different densities over an interval – the interval of precision depending upon the prevailing measuring technique [Popper 1965: 126].

So ΔP is somehow distributed to reveal the interval of precision and this in turn depends upon our technique of measurement. But as we have seen, we don't try to falsify a technique. So it looks as though we are testing (trying out) the instrument so as to provide us with information about its accuracy (interval of precision) as a measuring instrument. Popper also suggests that we should raise the degree of precision in our measurements as much as possible [Popper 1965: 124]. Presumably this is not just to do with determining an instrument's degree of accuracy, but trying to improve its accuracy as a measuring device. All this seems to depend upon the success of the applied sciences in light of informative testing.

Maybe Popper could argue that the methodology he describes is about eliminating sources of error in our measuring devices, and that the success of this error testing is a measure of a degree of corroboration for a theory of measurement.[26] Popper might then claim that there are an infinite number of problems that might affect an experiment. We cannot test for them all, so we are not justified in claiming that we are measuring accurately.[27] As a result, the decision as to what problems need to be tested for will involve guess-work and are thus risky. As noted, a falsificationist is quite comfortable with this result. Unfortunately, it also seems to turn applied science into a pure science. With so much that can go wrong in an experimental setting (or any setting where an instrument is being tested) it looks like our applied scientist takes some considerable risk when hypothesising about a source of an error (or a source of a possible error).

2.6 Designing an Informative Test

It might seem that the project, so far, has been largely negative. Although I have argued for the need for a test that might establish the applicability of a theory (the informative test) I have not provided much insight into how this might be achieved. In this section I briefly consider the work of Deborah Mayo who, using a broadly falsificationist methodology, has provided a way of assessing the degree of evidential support a test provides for a hypothesis. Although sympathetic to this project I think more is required. What is left out of Mayo's account will be explored further in the next chapter.

Mayo argues that science at the experimental level is a process of error elimination. Unlike Popper, however, she claims this process can result in a severe test that provides support for the hypothesis (H) under test. Confirmation, in this sense, amounts to finding a theory that survives the severest error test required. Hence:

[26] The relationship is far more complex than what is suggested here. In an appendix Popper develops his theory of corroboration and its relation to the kind of probability statements associated with error testing [Popper 1965: 406-19].
[27] There is some textual evidence that suggests he would take this approach [Popper 1965: 418-9].

> Evidence *e* should be taken as good grounds for *H* [the hypothesis] only to the extent that *H* has passed a *severe* test with *e* [Mayo 1997: 246].

It might seem that, given that there are many hypotheses that are consistent with the evidence, there is no way of deciding which of the alternatives are supported by the severe test (the under-determination problem). However, according to Mayo not all hypotheses consistent with the evidence are equally tested in a severe test. Once we get to the experimental context, using established background knowledge we can determine the relevant possible errors for the hypothesis under test, as well as finding a method for isolating and testing them. For example, Jean Perrin, when testing for Brownian motion, was able to group external factors that might feature as a possible source of motion. In this context the test need not be any more precise than identifying external errors as following a 'specified coordinated pattern' [Mayo 1997: 258]. Whatever the explanation is for Brownian motion, the fact that these kinds of external motion were eliminated as a source justifies the theory that there is an internal source of the random motion.

Mayo highlights the actual statistical tools used in experimental sciences used to assess the severity of a test. Here error is usually considered as the 'no effect' or null hypothesis (H_o) that might account for the evidence. If the difference between the observed mean result and the null hypothesis (H_o) mean result is two standard deviations, then we are justified in rejecting the H_o. A difference of two standard deviations represents a probability (P) of 0.03 that the actual test result is due to H_o. Importantly, Mayo argues that the rejection of H_o provides support for H in its passing of this severe test. 'Calculating severity for passing H is one minus the probability of such a passing result, when in fact the results are due to chance, i.e., when H_o is true.' [Mayo 1997: 257]. Mayo, it seems, is able to make this move because experimentalists are capable of isolating the relevant possible errors. There is not, as Popper thought, an infinite number of errors that need to be eliminated. By eliminating these errors (or testing for their likelihood) the test garners support for the hypothesis that accounts for the result. This measure of support is not absolute but relative to the test. Thus, different hypotheses, consistent with the evidence of the test, need not garner support from the test. According to Mayo, these alternatives may err in different ways, and thus are not severely tested. We can of course design a new severe test, the passing of which provides these alternatives with support. A hypothesis that suggests no new test, and is equally tested by the current severe test, is not, at the experimental level, an alternative.

Although Mayo's view is promising, Gregory Wheeler argues that, just by utilising Mayo's error statistic (*ES*) approach, we can always find an alternative hypothesis that is part of the model under test which has not been error tested. But now we have undermined the severity of the test, and thus the support the test provides for the model. Wheeler concludes: '[t]o the extent that *ES* solves Duhem's problem [the underdetermination problem], even partially, it does so by relying heavily on a rich body of knowledge that can't be accounted for by *ES* methods.' [Wheeler 2000: 418]. Yet we should note that Mayo does not deny the importance of background knowledge in solving problems, especially in what might count as a legitimate source of error. The *ES* approach is meant as a method to justify the decision to accept a hypothesis as a result of a severe test in the context of established background theory. While background theory is used to help identify potential sources of error, the *ES* approach provides a way of assessing or quantifying our confidence in the results of the experimental setup. Still, it is not clear how the error statistics approach might support the use of this background knowledge.

Mayo also claims that the ES approach can be quite informal. To explain this she uses an example, provided by Ian Hacking, involving inferences we might make when using a microscope. After describing two different physical processes to detect a body using a microscope, Hacking suggests that:

> It would be a preposterous coincidence if, time and again, two physical processes produced identical visual configurations which were, however, artefacts of the physical processes [Hacking 1983: 201].

Although this reasoning seems like an informal allocation of probability – a preposterous coincidence ($P<<0.03$) that the result were artefacts of the physical processes (H_o) – Mayo fails to note that Hacking immediately goes on to claim:

> Note that no one actually produces this 'argument from coincidence' in real life. One simply looks at the two (or preferably more) sets of micrographs and sees that the dense bodies occur in exactly the same place in each pair of micrographs. That settles the matter in a moment. [Hacking 1983: 201]

Hacking, it seems, is pointing to an inference that involves more than what Mayo is advocating. I am sure there are examples of informal error reasoning Mayo is alluding to. However, if Hacking is right, there is another important kind of 'inference from intervention'. The point to be drawn here is that the error statistics approach, although providing a way of measuring our confidence, does not explain the relationship between the evidence and the relevant experiment. And it is this sort of causal story that Hacking is attempting in his 'interventionist' approach to entity realism. Still, there does seem to be an important relationship between causal relations our instruments of measurement and detection inter into, the theories used to interpret such relations, and the testing in regards to this relationship. Mayo and Popper have provided considerable insight into what it is to test. How a theory might be legitimately applied so as to interpret, with some warrant, what is indicated by the functioning of an instrument is the central goal of the informative test. What constitutes an informative test has been left largely unanalysed, but at least its purpose should be clear. Making sense of such a test will be the subject of the next chapter.

2.7 Realism and the Informative Test

I have argued that we require an informative test in order to explain the success of the whole scientific enterprise. This seems a good result for those who hope for some secure theory with which we might interpret our instruments of measurement and detection. Using the informative test we are provided with information regarding the applicability of a theory. The information garnered from these tests allows us to apply the theory describing the empirical world in just those settings where the theory has been established as applicable. Moreover, it provides some confidence in the information we can expect from such an application. If this emphasis on the empirical utility of theory is interpreted as instrumentalism, so be it. The constructive empiricist I am trying to convince to be more realistic is hardly going to worry about this criticism. If I am successful in extending empiricism beyond the observable range, it hardly matters that some might interpret the theories used in this way are being treated instrumentally. However, I do need some stability in these ‘applied laws’. This chapter has argued that this stability is ensured, not because theories when considered as instruments are unfalsifiable, but because we have ways of testing instruments for applicability. I have paid particular attention to the sceptic of theoretical stability to show that something like an informative test is required. Still, there is some mystery that remains regarding the nature of information derivable from tests involving the detectable world. It is to this mystery that the next chapter turns. Fortunately, it turns out that this relationship is no more mysterious than, and analogous to, the link between perceptual judgements and the observable world. A relationship the constructive empiricist is prepared to commit to.

Chapter 3: Detection and Constructive Empiricism

3.1 Introduction

In the previous chapter, I argued that there are tests, or features of tests, that allow us to acquire information about the empirical world. To achieve this, the theories used to garner this information must be 'secure'. It was noted that theories, when considered as instruments, are not the kind of thing that can be true or false, and thus they can be considered unfalsifiable. However, the unfalsifiability of an instrument is somewhat beside the point. The important ability is our ability to test an instrument for its accuracy and applicability. Nor does the testing of instruments limit our ability to test a theory, though it does explain why we feel comfortable with the empirically adequate (applicable) parts of it. This should not be too controversial for the constructive empiricist, especially if we limit what is empirical to what is observable. In this chapter, I want to extend what it is to be empirical to include the detectable parts of the unobservable world. To do this, I hope to draw an analogy between what it is to observe and what it is to detect.[1]

This extension of empiricism into the detectable world I term *detectionism*, but there is no real need for the term if we accept that what is detectable is empirical. The important thing we need is some account of how we might gain epistemic access to the hidden but detectable parts of the world. As we will soon see, my analysis of detection is importantly related to operationalism. While operationalism has been the subject of considerable philosophical criticism, and although this criticism was seen to be fatal, I draw upon some of the key aspects of the approach that survived the attack. Properly developed, an operational method, coupled with our ability to test instruments of measurement or detection, explains how information about the detectable world may be obtained from what our instruments indicate. This will shed further light on what it is to perform an informative test and thus justify the results of such testing.

This account of detection by itself is unlikely to convince a constructive empiricist. It is arguable that, unlike what is observable, what is detectable is contestable. What we interpret when detecting depends on what theory the interpreter holds. Conversely, what is observable, according to the constructive empiricist, is a 'theory-independent question' [van Fraassen 1980: 57]. It is not clear what van Fraassen means by this term, but it is plausible that he is pointing out that facts regarding what is observable are not subject to change under various theoretical models, or not in any serious danger of being falsified by changes in theory. Unfortunately, van Fraassen does not provide a detailed account of the relationship between the observable world and our perceptual

[1] Godfrey-Smith [2003: 183-186], also suggests the use of detection as an epistemic criterion but does not go on to develop it.

judgements regarding it. Even so, he has provided some analysis, that, when coupled with the examples he uses, may provide the basis for a credible account of observation. To develop this account, I draw upon the work of Fred Dretske whose ideas on the subject seem consistent with van Fraassen's. Using this analysis, I will then argue that the relationship between the observable world and the facts about what is observed is analogous to the relationship between the detectable world and those claims concerning what is detected. Whatever it is that secures facts about the observable world, then the analogous features should do the same job for claims regarding the detectable world. The constructive empiricist may of course reject my analysis of observation, and thus its relation to detection. To guard against this, I also provide a simpler (although less informative) analogy that does not rely on the detail of the analysis. I consider responses to this extension of the empirical to include the detectable world in the last section.

3.2 Operationalism and Detectability

Influenced by the work of Einstein, Percy Bridgman, one of the founders of operationalism, tried to develop an approach to concept definition and introduction that accurately reflected scientific practice.[2] Bridgman was particularly impressed with Einstein's theory of special relativity. According to Bridgman, Einstein did not presuppose an ontological position on the concepts of *time*, *space* and *simultaneity*. Rather, by pointing to a series of physical operations, he provided what seemed to be an uncontroversial method for determining values for these concepts. If Newtonian theory was right, nothing interesting would come of this operational approach. However, Einstein famously conjectured that if we want to hold these definitions true in any moving frame (moving relative to another frame), there would be interesting implications when trying to reconcile the values measured from the perspective of each frame. For Bridgman (et al.), this indicated an important relationship between a concept and the physical method by which it was determined.[3] Rather than have a developed concept prior to determining what is measured, the priority was reversed: each kind of operation determined the meaning of each concept.

The strong implication of this was developed in Bridgman's early work. Here he suggested that different methods for measuring length, for example, were in fact defining different *types* of length. Consequently, the use of theodolites and survey chains measured a different type of length compared to those methods used to measure the wavelength of a gamma ray. It was a different type of length because the concept was nothing more than the operations used in evaluating it. Early on Bridgman offered the following strong thesis:

---

[2] Bridgman rejected the suggestion that he was founding a new philosophy of operationalism. Rather, he saw himself as highlighting how the meaning of a concept is established in the physical sciences.
[3] Einstein was later to shift away from this 'positivist' position towards a form of realism, though what this realism amounted to is contested [Fine 1984].

> [A] concept … involves as much as and nothing more than the set of operations by which … [it] is determined. In general we may mean by any concept nothing more than the set of operations; *the concept is synonymous with the corresponding set of operations*. [Italics original] [Bridgman 1927: 5]

Although sympathetic to one of the goals of operationalism – clarity of meaning – Robert Lindsay argued that it could not do the work it promised [Lindsay 1937]. Lindsay's central criticism was that concepts have an important role in unifying science, and Bridgman's operationalism was unable to explain this feature. Theoretical physics, for instance, has been successful in unifying various operational definitions under a single concept. When measuring pressure we can either use a U-tube, coil-pipe gauge, ionisation gauge or many other methods. Operationalism cannot explain what it is to unify using a concept, or whether or not we should.

Bridgman's response was to weaken operationalism. In a later paper, he claimed:

> We are not attempting to set up a theory of meaning and do not maintain that meaning involves nothing more than operations. We are dealing with a necessary as distinct from a sufficient characterization … [Bridgman 1938: 116]

In his later work, Bridgman also emphasised the pragmatic dimension to the meaning of a concept. We can, for practical purposes, use concepts to unify. Nevertheless, we should never forget that what may look like the same concept can, at a later date, be found to be made up of several distinct kinds. Bridgman explains that although the concept of length may involve two or more different operations, there is always the chance that, as improvements in the accuracy of our instruments develop, we will discover a difference. If there is a difference, the different operations cannot be defining the same concept. By 'forgetting' that the concepts currently unified may in fact be distinct, the unifying character of science risks an opportunity to advance.

According to this more nuanced view, theoretical concepts should be understood as having both a 'broad' and 'narrow' sense. The broad sense can include purely verbal or mental operations, 'pencil and paper' operations of mathematics, physical analogies or diagrams. Bridgman is not entirely clear here, but it seems that broad operations serve to unify. Several operations – the verbal, mental, pencil and paper kinds – can be used to define broadly the same concept. However, this unification only serves a useful though potentially deceptive categorisation. To be 'objective', nested within this broad set of operations there needs to be narrow definitions that involve physical operations. These various narrow interpretations make precise, for the scientist, the broad and potentially deceptive meaning.

This response, however, just seems to shift the problem. Sometimes an improvement in the accuracy of an instrument is for pragmatic reasons only. Arguably, it is the 'objectivity' of a broad concept that allows us to improve the accuracy of an instrument that measures something a concept refers to. If we could not hold the broad concept as 'objective' while improving the accuracy of an instrument, it is not clear how we should interpret the improved accuracy. The use of lasers to measure distances cannot improve on systematic errors in the old technology unless we assume that there is some objective concept *distance*, the property of which, some instruments can measure more accurately than others do. Hence, operationalism does not provide adequate principles for assessing when to prioritise concepts defined by individual physical operations over concepts used to unify. The question as to whether or not different physical operations define the same thing we might call the 'definition problem'. We will return to this problem shortly. In the meantime, there is a second and related problem.

Although the operational method professed to be an approach to concept definition, it should also be evident that the approach involves some ontological claims. The concept defined by virtue of the operational method has some relation to the property measured. Indeed, science is in the business of discovering properties, not just defining the theoretical concepts involved. Yet, if there is only one physical way of measuring the property involved, as the operational approach allows, then there seems no way of independently checking what is measured. If the property the laser-operation measures is not the same thing as the rod-operation measures (at least potentially), or any other *x*-operation that measures an *x*-distance, how might we check what is being measured by each operation? Let us call this the 'measurement problem'. We will have to be careful with this problem. Solving it might give the impression that we are begging the question against the constructive empiricist.

There is a third type of problem for the operational approach. One of the advantages of the approach, it seemed, was that it allowed us to refer to some physical process as a basis for concept definition. Unfortunately, as Popper [1965: 440n.] notes, this does not avoid a circularity problem. For instance, an operational definition of length that involves solid rods will also require a correction for temperature. Yet the 'usual' definition of temperature involves a measurement of length. I think Popper is right about this, and thus we will require more than operationalism can provide, but I also think that matters are more complicated than suggested by this kind of example.[4] For a start, it is not clear how vicious the circle is. Consider the theory defining a force between charged particles.

1. $F = q_1 q_2 / r^2 4\pi\varepsilon_o$

[4] Popper also suggests that there is some sort of regress of testing. This problem, or something like it, is developed further by Fredrick Suppe [1972].

The operations of some sophisticated torsion balance involving charged spheres (along the lines developed by Coulomb) might do the job of defining *F* according to this theory. Yet in order to understand this theory we need to know what the other terms in the theory mean. To do this, we might operationally define them. So for example, *q* in a different context might be defined using the following theory:

2. $q = F/v \otimes B$

The set-up (operation) of a cyclotron could be used to measure *q* and thus operationally define it. Unfortunately, as Popper would have pointed out, *F* also appears in the operational definition of charge. Even so, maybe we can define *F* in another context without reference to charge, say:

3. $F = MA$

Now it is true that this definition will require an operational setting, and as a result we have introduced more concepts that require defining, but at least it seems to avoid the circularity problem identified by Popper. This solution also highlights the related measurement and definition problems. We will need some way of claiming that *F* mentioned in (*3*) and defined in terms of its operations is the same concept as *F* mentioned in (*1*). Furthermore, we need some justification that both sets of operations can measure the same property, if in fact anything is being measured. Popper might have also pointed out that the solution relies on the addition of new concepts. If we insist on avoiding circular definitions, it looks like we end up in a regress involving the addition of new concepts.

Still, maybe we can play some of these problems off against each other. Something like this seems to be behind the approach offered by George Schlesinger [1958]. Schlesinger makes the point that there is an important aspect to Bridgman's operationalism that he thinks establishes a 'differential meaning'. A differential meaning is established in virtue of two or more independent operations that define the same theoretical concept. The different operations required by two different theories (in which the same term appears), at least initially, can identify two different concepts. However, if it so turns out that we can measure the same values for the theoretical term using different operations, then the term that features in both theories has a 'differential meaning'. For example, Hempel (et al.) had noticed that operationalism could define empirical concepts that lack 'theoretical import'. Consider the concept of 'hage' – the product of height and age [Hempel 1952: 46]. Although this concept can be operationally defined, there is no theory that might relate it to other concepts. From the differential perspective, *hage* lacks a differential meaning because there are no other independent operational definitions that might articulate the concept. In contrast, the concept *stress* has a differential meaning. Stress can be defined in terms of those operations that determine the force acting on a unit area, or by those operations that determine a rate of deformation (strain) relative to a modulus of elasticity.[5] Because the two *independent* operational definitions can be used to measure the same property – stress – the concept has a differential meaning.[6]

Schlesinger goes on to suggest that establishing a differential meaning provides us with 'the strongest evidence that we are dealing with a 'real' property transcending the methods defining it' [Schlesinger 1958: 305]. Although I think the differential approach is important, for the purposes of my argument, the analogy between observation and detection will establish what we can commit to as real. Nevertheless, the solution to the measurement problem that the differential approach provides has something to do with *what* is measured. It seems to rely on playing off the measurement problem against the definition problem. The different theories involved must evaluate (provide a value for) what it is that is detected or measured. If the evaluation cannot be made equivalent, it looks as though we are measuring or detecting different things. If they can be interpreted as equivalent, this warrants the belief that we are measuring or detecting the same thing. If the latter is the case, the theories, in combination, differentially define whatever it is that is measured by the two operations.

Schlesinger's insight is important for the type of realism I am developing, but it is still a bit mysterious how the evaluation is supposed to work. Hempel provides a similar analysis of this kind of operational approach, but draws a different conclusion. According to Hempel, the equivalence of the two operational definitions implies a law that has to be tested empirically; the equivalence cannot be assumed. If this is right we might be able to embrace the circularity problem to allow us to test the 'definitions'.

Before we go on to consider this idea in more detail, we might pause to consider what I think the intuition is behind the differential approach. The idea seems to be that there is some important relationship between making some measurement, and our ability to interpret what the result of such a measurement is. Presumably, the similarity in measurements, if differentially established, has something to do with the functioning of an instrument or operation of measurement. The regular positioning of the dial on a gauge indicates something (or not) regardless of our (mis)interpretation of what has been measured or detected. For reasons that will become evident in the next section, I

[5] Stress can be represented in one theory as $\sigma_x = F_x/a_{yz}$ where the suffixes represent the relevant Cartesian axis. F_x is force along the *x*-axis; a_{yz} is the surface area in the *yz* plane upon which the force acts. The relationships can also be stated as the generalised Hooke's Law for a three dimensional body in a Cartesian space: (σ) stress; (ε) strain; (E) Young's modulus; (υ) Poisson's Ratio (Torsion is ignored here for simplicity) $\varepsilon_x = 1/E\ [\ \sigma_x - \upsilon\sigma_y - \upsilon\sigma_z]$, $\varepsilon_y = 1/E\ [\ \sigma_y - \upsilon\sigma_x - \upsilon\sigma_z]$, $\varepsilon_z = 1/E\ [\ \sigma_z - \upsilon\sigma_x - \upsilon\sigma_y]$.

[6] It could be argued that *hage* also has a differential meaning because the concepts *height* and *age* have differential meanings. However, this just indicates that the meaning of this concept is derived from differentially established concepts. Whether or not *hage* has a non-derived differential meaning remains to be established.

term this indicative nature of an instrument, ‘non-epistemic’ measurement or detection. This non-epistemic aspect of measurement and detection allows us to interpret what is detected and measured, but not in a way that makes it entirely arbitrary. Fred Dretske [esp. 1988] has made much of the analysis of an instrument’s capacity to indicate, and I will be drawing upon some of this analysis later. The important point being made is that if we have two independent instruments or operations that indicate the same feature, this provides us with the opportunity to check what is being indicated. If we only have one instrument that might indicate, then it seems there is no way of checking what the instrument is indicating.[7]

Returning to the differential approach, as Hempel noted, empirical laws can be found to relate two independent operational definitions. This might prove Popper’s point that these definitions are circular. However, if Hempel is right about the testability of the implied ‘empirical law’, the circularity does not seem so vicious. Importantly, as Hempel also notes, if testing is important to the analysis of concepts using the operational approach, the approach can no longer be about definition. Rather than see this as a problem, I think we should embrace this result. The operationalism of the differential approach is more appropriately construed as a methodology for assessing our ability to *detect* or *measure* the properties to which theoretical concepts refer. Having made these claims, we might now see how it works with some examples.

Donald Gillies, who also criticises a strong version of operationalism, provides a viable account of how to test theories from an operational perspective [Gillies 1972]. Gillies utilises Newton’s introduction of the concepts of *force* and *mass* as a historical archetype of the operational method. Given that these concepts were new, we might ask, how was it that Newton provided ‘empirical significance’ or ‘meaning’ to them? More generally, what is it to provide empirical significance? For Gillies:

> A concept has ‘empirical meaning (or significance)’ if we can assign numerical values to particular instances of it – if we can, in effect, measure it under certain circumstances. [Gillies 1972: 8]

This formulation, as it stands, allows too much to be empirically significant. There is more to the empirical meaning of a concept than just assigning values to it. We could, it seems, give a value to any fictitious concept. If this form of empirical significance is not further qualified, a thermometer, for example, could indicate the number of fairies in a liquid. Gillies might contend that, operationally, a thermometer *can* measure the number of fairies in a liquid. The number of fairies would just mean what the thermometer measured. We must not presuppose what word ‘fairies’ means in this setting. However, if we haven’t already presupposed what a thermometer measures, it is not clear that we

[7] The term ‘instrument’ should be considered quite broadly to include what natural phenomena might indicate.

can appreciate what it is measuring. If we have no idea what an instrument is measuring, even if we can assign a value, what is measured will lack empirical significance. So it looks as though we need to presuppose that the concept is empirically significant in order to provide it with empirical significance. As we will see, Gillies' account does try to attend to this problem. Even though he is only partially successful in avoiding this issue, Gillies does provide a way of testing operationally 'defined' theories.

As noted above, often the same theoretical terms appear in different theories. Even if the concepts are new this feature might allow us to test one theory against another and thus in some way introduce new concepts in a non-circular way. Simplifying the insightful reconstruction that Gillies provides, we can consider two of Newton's equations that involve the concepts of *force* and *mass*.

4. $F_m = MA$

5. $F_p = \gamma (M_1 M_2 / r^2)$

Combining these equations we can derive a third equation describing the geometry of the orbit that is the result of the mass-force relation.[8]

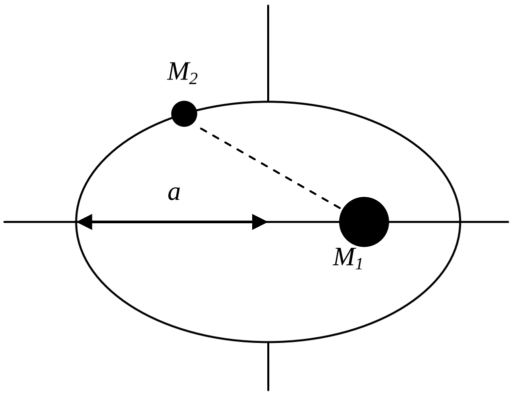

6. $a^3 T^2 = \gamma (M_1 + M_2)/4\pi^2$

In the above formulation (6), T is the period taken for M_2 to trace the ellipse about M_1, with a being the major semi-axis of the ellipse as shown. We still need some assumptions to assign values to the mass concepts in the equation and thus test (6) with respect to the astronomical data. According to Gillies, if one of the masses M_1 is the sun, we can infer that its mass is very much greater than M_2 – the orbiting planet. This reduces (6) to:

7. $a^3 T^2 \approx \gamma (M_{sun})/4\pi^2$

[8] Not quite: (*4*) and (*5*) presented above are in their scalar form; their vector analogues are required to perform the derivation. Equation (*6*) also assumes a fixed point mass M_1.

Although it is a matter of convention what value we assign M_{sun}, we would not be able to find a constant value for γ for planets affected by the sun if the above account was wrong.

We can see this example depends upon something like the differential approach. Each theory is initially understood in terms of its own physical operations. One operational setting uses equation (*4*) and those measurements in which a force is applied to a free mass, which duly accelerates. The second setting uses the equation (*5*) to measure a force generated by gravity proportional to the masses involved. We can, using equation (*7*) (Hempel's bridging law), test what we interpret is indicated in the various operational settings. If successful, this (differential) approach justifies the belief that what the concepts refer to is measurable and thus differentially established. If on the other hand what is measured is inconsistent with what is outlined in (*7*), we are justified in falsifying one or all of the theories involved.

Although this approach looks promising, it is not sufficient for establishing empirical significance. To see this we need to consider the assumptions required for the test: assumptions such as, that the sun's mass is much greater than the planets that orbit around it. Of these assumptions Gillies asks:

> Do we need an operational definition of mass at this point? Not at all! We test our theory involving masses by making the qualitative physical assumption that one mass is very much greater than another. Moreover *this qualitative assumption is justified by a crude (or intuitive) notion of mass*. If we think of mass as a 'quantity of matter', then observing that the sun is very much larger than the planets and making the reasonable postulate that the density of matter is at least comparable to that of the matter in the planets we obtain $m_s >> m_p$ [Gillies 1972: 13; my italics].

The important point to note here is that the test is not possible without an intuitive notion of mass – the *new* concept, for which we are trying to provide cognitive significance. Even though I think Gillies has shown us how we might test two theories operationally construed, empirical significance of the new concepts requires some intuitive basis.

A similar method is applicable to instruments rather than natural phenomena. The relationship between theory and instrument does not differ in any substantial way from the method outlined above. Ensuring the correct measurement is a process of testing at least two theories in which the same theoretical terms appear against the instruments over which they range. Here we can consider how the concept of *temperature* might be differentially established and thus measured. Initially, an operational setting correlates a temperature change with, say, the displacement of certain materials – mercury or alcohol. Using some established theories about the properties of other materials, an instrument can be designed and manufactured: more or less a thermometer.[9] It is true that we have had to assume a concept of *temperature* and its relation to a concept of *displacement* to formulate this operational setting, but we need not judge what the instrument measures just yet. Importantly there is a second theory that suggests that temperature has a fixed value at the freezing and boiling points of distilled water. Once again, the term we are trying to define appears in this operational setting. However, if we combine the two settings, testing a thermometer immersed in distilled water as it ranges between boiling and freezing points, we can determine if temperature is measurable and thus differentially established. If we were wrong about the relationship between the two operational settings, measuring temperature would not be possible in any reliable way. It is reasonable to suppose that none of this is possible without an intuitive concept of *temperature* and *displacement* to start with.

Gillies' account has shown us how to test one operationally defined theory against another. We have also discovered that we need some intuitive concepts. These intuitive concepts are importantly involved in the development of the theory. Yet these intuitive concepts themselves do not live in a theoretical vacuum. For instance, there seems to be some 'theory' that suggests that whatever the concept *mass* means, it is in some sense proportional to volume. At this stage we need not make too much of this point. However, I will have a little more to say on this matter when we come to consider the nature of observation.

Finally, we should note that some of the interesting features of detection I have uncovered in this analysis also inform Ian Hacking's work. Hacking [1985: ch. 11] also makes use of the idea that two independent modes of detection can settle the matter as to what counts as being detected. Hacking's entity realism, however, tries to do this without any commitment to a theory. As we saw in chapter 1 this exposes his account to a problem. The fact that we can say very little about what is detected left us with a rather empty sort of realism. In contrast, the differential approach looks as though it may be able to provide some content to what is indicated by the correct use of our instruments of detection. Of course, in order to convince the constructive empiricist the analogy between detection and observation still needs to be made.

3.3 The Analogy: From Observation to Detection

Having analysed the conditions under which we might detect, we are now in a position to make the analogy with what it is to observe for a constructive empiricist. If the two

[9] It is obviously far more complex than what is suggested here; though, I do not think the complexity undermines the method.

achievements are not relevantly different, then there is no additional risk in extending what can be known empirically. Before attending in detail to the analogy, I might first highlight what I hope to show is relevantly similar between observation and detection. From the above analysis, we can distil five important conditions required for successful detection:

(i) In order to detect some feature we require at least two different operations or modes of detection capable of detecting or measuring the same feature.

(ii) The operation of an instrument of detection or measurement, when detecting or measuring, must indicate, in some non-epistemic way, some feature such that we can infer (maybe mistakenly) what we have detected or measured.

(iii) A theory is required to interpret what is indicated by each method of detection or measurement.

(iv) When interpreting different modes of detection or measurement, if each interpretation is not consistent, we have reason to question just what it is that has been detected or measured. If the measurements are consistent, we can accept what the instruments are indicating.

(v) Concepts, related by an overarching theory, are required in order to begin the process of interpreting just what we believe we have detected or measured.

Ideally, I would like to show that something analogous to these five features is evident in the account of observation that the constructive empiricist has provided. Unfortunately, van Fraassen does not provide a detailed account of what it is to observe. So part of my project will be to show that these conditions are required of an account of observation that a constructive empiricist ought to commit to. I aim to show that something like these five features allows the constructive empiricist to commit to simple perceptual judgements regarding what is observable. If these conditions allow us to commit to perceptual judgements regarding the observable world, then the analogous conditions allow us to commit to what we judge is detected of the detectable world.

We might start our analysis by noting that van Fraassen defends constructive empiricism against arguments that deny that there is a distinction between what is observable and what is unobservable. As we will soon see, he needs more than a defence of this distinction alone. In order for a theory to be empirically adequate, not only does there need to be something observable, the theory must *also* accurately describe what is observable in an empirically adequate way. We can see this double feature in the following:

> [A]… theory draws a picture of the world. But science itself designates certain areas in this picture as observable.

This represents the first distinction between what counts as observable in the picture a theory draws with the rest of the theory. Van Fraassen immediately goes on to point out that:

> The scientist, in accepting the theory, is asserting the picture to be accurate in those areas. This is, according to the anti-realist, the only virtue claimed which concerns the relation of theory to world alone [van Fraassen 1980: 57].

Science is responsible for designating which part of a theoretical picture is observable (and which part is unobservable); of the part of the picture that is observable, the theory's description of this area has to be accurate. Van Fraassen also notes that what he counts as observable in this picture is a theory-independent fact.

> To accept the theory involves no more belief, therefore, than that what it says about observable phenomena is correct... This might produce a vicious circle if what is observable were itself not a simple fact disclosed by theory, but rather theory-relative or theory-dependent. It will already be clear that I deny this; I regard what is observable as a *theory-independent* question [van Fraassen 1980: 57; my italics].

Van Fraassen immediately goes on to qualify this remark. Although the facts regarding what is observable for us humans are, in some sense, theory-dependent, 'there is not the sort of theory-dependence or relativity that could cause a logical catastrophe here' [van Fraassen 1980: 58].[10] It seems then, that what this independence amounts to is a freedom, on the part of science, to investigate and formulate theories about what is observable, without destroying the distinction between what is observable and what is unobservable.[11]

Apart from the distinction between the observable and unobservable, there must be some important epistemic relationship between observable objects and the claims regarding such objects. Unfortunately, we are not provided much guidance on what should count as a correct description of the observable phenomena. This, according to van Fraassen, 'will be described in detail in the final physics and biology' [van Fraassen

[10] The sense in which what observable for humans can be 'theory-dependent' but avoid a vicious circle is developed further in van Fraassen [1992].

[11] It seems we still need another kind of 'theory-independent' fact other than those facts that describe what is (and what is not) observable. In order for a theory to describe accurately the empirical world, we need theory-independent description of the observable *phenomena*. It is not enough for all theories to agree as to what is, in fact, observable, but then *radically* disagree in the description of what is observable. If there is a radical difference in what is observed of the observable, then it is hard to argue that various theories can accurately describe the empirical world. So, apart from the facts that describe something as observable, there must be 'theory-independent' facts describing something of the observable phenomena. Mark Wilson [1985] makes use of a similar point in his criticism of constructive empiricism.

1980: 17]. Yet surely we can't wait for the results of these fields of science? Elsewhere van Fraassen claims:

> Our judgements of empirical adequacy of theories will of course vary: but whether those theories are empirically adequate – just like whether or not they are true – is a characteristic which they do not lose when we begin to think differently [van Fraassen 1992: 20].

This 'characteristic' has something to do with the observable world, not our opinions about what is observable. However, just saying this hardly explains the role the observable world has in securing some empirical adequacy for a theory. Nor do the illustrations that van Fraassen provides clarify this relationship. For instance, when articulating what empirical adequacy means in Newton's theory we are told that 'his [Newton's] theory has some model such that *all actual appearances are identifiable with (isomorphic to) motion* in that model.' The term 'appearances' here refers to ''apparent motions' [that] form relational structures defined by measuring relative distances, time intervals and angles of separation' [van Fraassen 1980: 45]. Unfortunately, it remains unclear how we might measure a relative distance *qua* relative distance, or a time interval, or an angle of separation. More plausibly, what we measure, if we can, is the relative distance or angle of separation between two identifiable – that is observable – points, objects, edges or whatever (and something analogous for a time interval). Presumably, this is how van Fraassen hopes to construe matters. But given that what counts as observable in this context is not mentioned, it is not clear how it might be related to the data of measurement.[12]

Elsewhere, when outlining the limits to empirical description we are told that we can, according to classical mechanics, 'designate as basic observables all quantities that are a function of time and position alone' [van Fraassen 1980: 59-60]. Excluded from the empirical data is a range of other 'quantities' such as mass, force, momentum and kinetic energy. Various theories can then quantify these features in their models so that they are at least consistent with the observable quantities. Yet the term 'basic observable' here seems to have a different meaning to what is humanly observable. It is not at all clear what might be involved when observing (unaided human perceiving) a quantity that is a function of time and position alone. In fact, a charged particle traversing a cloud chamber can be quantified as a function of time and position alone – van Fraassen's archetype of an unobservable entity. So it must be the case that the term 'basic observables' has some important relation to what is observable. Now it is true that the example is meant to make precise the point that, a 'physical theory cannot be translated, without remainder, into a body of statements that describe only what the observable phenomena are like' [van Fraassen 1980: 59]. Even so, we are left

[12] The potential complexity of this problem cannot be overstated. For some analysis of this complexity, see Laymon [1984].

wondering how this example of a description of what the observable phenomena are like is related to what is observable, and thus what it is to judge the theory for its empirical adequacy.

To disarm these concerns, presumably we are meant to keep in mind the pedestrian examples like observing a tennis ball or a car crash that van Fraassen has previously provided.

> Suppose one of the Stone Age people recently found in the Philippines is shown a tennis ball or a car crash. From his behaviour, we see that he has noticed them: for example, he picks up the ball and throws it. But he has not seen *that* it is a tennis ball, or *that* it is a car crash, for he does not even have those concepts ... To say that he does not see the same things and events we do, however, is just silly ... [van Fraassen 1980: 15].

Different theories can describe a tennis ball, but no theory can deny that the tennis ball is observable under each theoretical description. Van Fraassen, it seems, sees the confusion here arising out of a category mistake we might make when comparing what is observable with what is theoretical. The term 'theoretical' applies to terms or concepts, while term 'observable' applies to objects or entities. Be this as it may, it seems to me that there is more to this example than meets the eye. Like van Fraassen, I too find the above account convincing. The question I want to ask is, what makes it so persuasive?

There is a way of justifying the belief that we all see the same thing when shown a tennis ball that is consistent with constructive empiricism. In van Fraassen's example, the people from the different cultures are capable of picking the tennis ball up and throwing it. From this behaviour (ability), we can infer that they (or ourselves) are observing a tennis ball: even if the observer does not appreciate that it is a tennis ball. Nevertheless, we might well ask, why should this behaviour (ability) clinch the matter? Presumably, it is because we can appreciate that the tennis ball is the kind of thing that can be picked up and thrown. Compare this with a car crash, van Fraassen's other example of an observable. A car crash is not the kind of thing that can be thrown. There is obviously a different set of behaviours relevant to this observable (event). If a person tried to throw a car crash as if it were a tennis ball, we should infer that this person is severely mistaken about what is being observed.

Normal perceivers, then, see that a tennis ball is an object of the right size and weight to be thrown, and that trying to throw a car crash is to make a mistake. Importantly, we need not throw a tennis ball to establish that it is observable. We could just pick it up and squeeze it to see how if deforms, or test it for some other properties. By these actions we can know that the perceiver has formed some empirical beliefs. The beliefs that: firstly, the tennis ball is one of those things that is observable; and secondly, that the tennis ball has at least – how ever you want to describe them – those empirical properties required for observing, throwing, squeezing and so on.

Occasionally we are mistaken about what we observe. Something may only appear to be a hand-sized spherical object, maybe because of some trick of poor lighting. How is it that we can infer that we have been mistaken? Well, an important part of the inference could involve evidence from other modes of sensing, for instance from grasping and manipulating the said object. If, despite what appears (optically) to be the case, an object behaves not at all like a hand-sized spherical object as we manipulate it, we ought to call into question the evidence of our senses. Equally, if the two modes of perceiving are consistent in what they indicate we are justified in believing we have observed correctly.[13]

If this analysis is right, observing is not (just) about appearances (perceptual or tactile). It is also an epistemic achievement regarding what beliefs are justified, sometimes in spite of appearances. As van Fraassen notes, 'a complete epistemology must carefully investigate the conditions of rationality for acceptance of conclusions that go beyond one's evidence' [van Fraassen 1980: 72]. Nor, it seems, should we limit this epistemic achievement to ways of appearing. Following Dewey, an epistemology can draw inferences from ways of doing and manipulating. It seems that appearances cannot be manipulated in the way that real features of the world can. Unlike throwing tennis balls, we cannot manipulate, in the relevant way, what only appears to be a coloured arch in the sky, a rainbow. Even if this is going too far for a constructive empiricist, it seems reasonable to suggest that what helps in establishing the facts regarding what is observed, as opposed to mere appearance, depends upon what is indicated by different modes of perceiving. If this is right, the results of independent modes of perceiving is directly analogous to condition (i) identified of our analysis of detection. If this feature is important in establishing what counts as observed in the observable world, there is no reason why it cannot help in establishing what we detect of the detectable world.

Given what has been said so far, the issue now is, can we still maintain the distinction between observing *that* something is the case and just observing? As we have seen, van Fraassen hopes to defend just this distinction. Yet it is also evident that what counts as observable now depends on judging the relationship between certain observations (actions and behaviours). I think the distinction can be maintained and I will draw upon the work of Fred Dretske's [1969] to show how. This analysis then sheds further light on what else is required to commit to beliefs regarding what is observable.[14]

[13] Obviously there would be further defeating reasons to be considered in a detailed account of an epistemology based on this principle. However, I think the principle itself is plausible enough.

[14] The development of the analogy also draws upon ideas presented in Dretske [1981] and Dretske [1988].

Dretske's understanding of observation is that it has both epistemic and non-epistemic aspects. He terms this non-epistemic aspect 'seeing$_n$'. The non-epistemic aspect of seeing is what is common to everyone. True, there are scenes where we can view a 'change' in what is seen depending on what we believe or know; the usual suspects include duck/rabbits, Necker cubes and hidden faces. Nonetheless, there is still some content for all to see$_n$ regardless of what we see something as. All people see$_n$ the world the same way, and seeing$_n$ is just as rich for all people. On the other hand, epistemic seeing, characterised as 'seeing that …' are those perceptual judgements we form about the world that we all see$_n$ in the same way.

Non-epistemic seeing provides further insight into van Fraassen's point that the term 'theory' in 'theory-laden observation' is applicable to linguistic terms, whereas the term 'observation' is a perceptual activity or a process applicable to objects or entities. It seems we are making a category mistake when claiming that an observation, rather than an observation statement, is theory-laden [van Fraassen 1980: 14].[15] Even if a theory we happen to believe causes (some) changes in what we see$_n$, it would be wrong to claim that because of this, observation is theory-laden. The cracking of a whip might cause the horses to run, but this does not make whip cracking part of their nature. This seems to be the mistake that Russell Hanson makes in his argument for his thesis that seeing is (entirely?) theory-laden [Hanson 1972: esp. ch. 1]. Merely pointing to the fact that what we believe may change some aspect of what we see is not sufficient to show that what we see, rather than what we believe we see, is theoretical. These matters are complex, but even if seeing is partially theoretical, then it is reasonable to ask, which part or to what degree is seeing theoretical? Of the part that is not theoretical (or the degree to which seeing is not theoretical) what might we claim of it? Whatever the case is, Hanson's analysis cannot imply that there is nothing in common between different observers of the same thing. Without something in common in what is seen, it would not be possible to propose the thesis, as Hanson does, that Kepler and Tycho saw the same thing – the sun – differently.[16]

Returning then to our tennis ball example, we might consider it again using Dretske's concept of non-epistemic perceiving. It seems reasonable to suggest that some tactile feeling$_n$ (what we feel$_n$ of a tennis ball when throwing it) is importantly related to some optical seeing$_n$ (what we see$_n$ of a tennis when viewing it). Of course when viewing a tennis ball we also view a whole scene, and when throwing a tennis ball we are also doing many other things. Putting aside this complication, there is some relationship between these two non-epistemic modes of perception that allows us to infer that we perceive the same thing. At the least, what these two sensory modes indicate has to be

[15] Van Fraassen, I think, views concepts as linguistic entities, or only expressible linguistically. For an alternative view, see Alva Noë [2004].

[16] In fact, in places, Hanson seems to acknowledge the precise distinction that Dretske and van Fraassen also make '[t]he gap between pictures and language locates the logical function of 'seeing that'. For vision is essentially pictorial, knowledge linguistic' [Hanson 1972: 25].

consistent such that we can judge that we have perceived correctly, or from someone's behaviour, that they see_n the same thing that we do when shown a tennis ball. But judging something as consistent (or not) is an epistemic achievement. Given that $seeing_n$ is non-epistemic, we need something else for us to form a belief based on what the different modes of non-epistemic perceiving, in conjunction, indicate. The most plausible contenders for this something else are concepts and maybe some theory that relates them.[17]

It should be clear that the non-epistemic seeing is directly analogous to the second condition (ii) required for detection. Non-epistemic seeing is, at least in part, what a mode of perception indicates. Also, if we are going to make these judgements about what each mode of non-epistemic perception indicates, we will require, in some sense, a theory to interpret what is indicated by our senses, condition (iii). Condition (iv) will be required if we are to judge whether or not what our senses indicate is consistent. Like observation, detection too required some non-epistemic identifiable feature such that we can infer (maybe mistakenly) from what the instrument or operation indicates we have detected or measured. If what is indicated is inconsistent, we should question our interpretation. If consistent, we have warrant for our judgement as to what is indicated by the use of the instrument. Like non-epistemic perceiving, non-epistemic detecting can inform us about the detectable world.

What role might the remaining condition (v) have in an account of observation? The nature of a concept is a much-contested topic, thus I am wary of wading too far into what looks like deep water. Even so, van Fraassen himself admits that perceptual judgements involve concepts, and are to some extent theory-laden. The question is, what kind of theory is required for those perceptual judgements that determine what is observable? To make a start on an answer, we might notice that different concepts can be importantly related. For instance, whatever the empirical concept *red* means or refers to, it has a certain important relation to another concept, *spatial extension*. Having a colour means having spatial extension. In fact, a colour without spatial extension seems empirically meaningless. Yet, even if a perceiver could correctly identify colours and

[17] The philosophy of perception is a complex and contested arena. I have chosen Dretske's account as I think his theory consistent with the small amount that van Fraassen says on the subject. However, I should not be read as endorsing Dretske's entire account. As should be evident, although I think Dretske is right to note a non-epistemic aspect to perception, I also think concepts are required for a complete account. Although I will not defend this position further, there are independent arguments for it; see for instance Maund [2003: esp. 33-40]. Maybe I am wrong about this. If I am wrong, then we still require some explanation for why some aspects of observation are shared by similar observers observing the same thing. Jerry Fodor [1984] has argued, for instance, that a conceptual account of perception can achieve this by positing a 'modular' aspect to how perception functions in a brain. Some concepts used for perceptual inference are fixed, some flexible. But this modular aspect can't be the entire story. If we are to be realists about observables in the world, then there has to be some relevant (causal) relationship between these observables and our fixed perceptual judgements about them. This relevant relationship has to determine, in alignment with the fixed concepts used, the fixed perceptual inferences made. Although I have not made use of this alternative account of observation, properly construed, I do not believe it threatens the analogy I am drawing.

what counts as spatial extension, the important relationship between the two concepts may still need explaining. In some sense then, the relationship between the concepts is theoretical.[18]

Extending this idea to other concepts, we can notice that having a mass means having a located volume. Also, a wave, it might seem (pre-scientifically), requires a material medium, while all materials have a temperature. Whatever the concept *force* is, it involves, further, the concepts of *location, direction*, *pressure* and *area*, as well as an entity to be forced. Maybe, we (humans) even have some intuitive understanding of a conservation principle; something cannot be made from nothing. Whatever the concept *material* means, it involves further not disappearing without a cause. Perhaps the relationship between various concepts is learnt; maybe some of this understanding is innate; maybe the relationship is independent of anyone understanding it at all. Whatever the case, the theories that relate empirical concepts explain, to some extent, our 'innate flair' (following Quine [1969]) for understanding the world in terms of natural kinds.

It is the nature of empirical concepts and the theories that relate them, I suggest, that is important when explaining why throwing and viewing are importantly related, and why, when considered in combination, these achievements allow us to commit to the observable tennis ball. As we noted above, what made a tennis ball observable involved a range of human modes of perceiving: e.g., seeing (optically) a hand sized weighted sphere and throwing (tactile perceiving) a hand sized weighted sphere. The empirical concepts involved (*hand-sized*, *weighted*, *sphere*), coupled with some theory relating these concepts, in this case to *throw-ability*, justify the inferences made and articulated as observable facts. These inferences are based on interpreting what a range of non-epistemic perception (and action) indicates. We often make mistakes about how concepts are related. These mistakes are revealed to us as we negotiate a world that tests our abilities, and thus their relation to the theories and concepts we use to interpret such results.[19] Yet, if we all (the epistemic community) have access to some similarly constituted empirical concepts and a theory that relates them, we can explain why different people see and do something similar by virtue of what their perceptions indicate. It seems as a minimum then, in order to interpret what is indicated by our senses and actions we require something like condition (v).

Still, it might be argued that I am attributing to all humans throughout history (not to mention other sentient animals) too much theoretical knowledge. Did Aristotle really

[18] It could be maintained that it is part of the concept *colour* that it has a spatial extension. I am sympathetic to this idea, but I think it is consistent with the claim that concepts have a theoretical component. Whatever the final analysis is, I think the distinction between concept and theory is worthwhile keeping.

[19] I am not suggesting that concepts are based entirely on these kinds of practical abilities, although Alva Noë [2004] has defended this idea.

understand what constituted the concept of *force*? Yes and no. Both Aristotle and Newton shared an *implicit* and basic understanding of the concept of *force* (as we all do). Newton made a theory involving the concepts *force* and *mass* explicit by evaluating it using something like the differential method. We can of course relinquish (with some difficulty) the basic understanding (theory) that relates concepts; we now have mass without a volume, spatiotemporally non-located particles, and effects without cause, but the decision to modify these important relations between concepts is not taken lightly. More importantly, even if we do find an exception to the way we usually understand a concept, this need not affect the clear settings where the application of the concept is well understood. Someone who believes, after observing a storeroom packed to the rafters, that it contains almost nothing, would be wrong. If what he sees_n is what I see_n – a store room packed to the rafters – the problem might just be that we are using different concepts and theory to interpret what our common seeing_n indicates. Maybe, once qualified with what we have detected of the molecular world, he can, in that context, be correct. The theory used to judge what we see_n is relative to a context, yet we do not have to worry about the threat of relativism. Testing the applicability of a theory is about determining the appropriate contexts for its use.

This account of the relationship between theory, concepts and perception is still a bit vague, yet I think enough has been said to see its analogous relationship to detection. In order to make various inferences about what our various ways of non-epistemic perceiving indicate we require empirical concepts and an overarching theory that relates them. Without something like conditions (i) to (v) we could not compare what our various modes of sensing indicated. We would not be able to judge our own perception or the perceptual claims of others. Importantly, if satisfying these five conditions established facts about the observable world, then the same feature can establish facts about the detectable world.

This account of observation has turned out to be quite complex. Maybe this is as you would expect, given the scope of the project. Maybe many aspects of this analysis van Fraassen would reject. Nevertheless, there would still be good reasons for him to retain at least the non-epistemic aspect to perceiving. Non-epistemic perceiving explains the precise distinction the constructive empiricist wishes to maintain between what is observable and what is theoretical. Yet even on this minimal construal, a simple analogy between perceiving_n and detecting_n will suffice to establish a realistic attitude towards what is detectable. If we consider the sensory apparatus of our bodies as instruments of detection [van Fraassen 1980: 58-9, Monton and van Fraassen 2003: §4.2], perceiving_n is just a special case of detecting_n. If this is right, then there is no reason not to extend what ‘empirical’ means to include all that is detectable. If it is possible to commit to the observable world with a sparse account of perception, then surely it is reasonable to commit to the detectable world on a similar basis. What is detected_n by an instrument in a non-epistemic way is just as theory-independent as what is detected_n by the instrument that happens to be the apparatus of human perception. The importance a constructive empiricist might place on the human perceptual system, rather than some other instrument, has to be justified rather than stipulated. And given that, when compared with the human perceptual apparatus, calipers, theodolites, thermometers and atomic clocks are all better at measuring accurately, and thus better instruments for assessing empirical adequacy, the special status given to the human perceptual system seems unwarranted.

It might be objected that humans have to *have* a concept in order that they might observe, while the relationship between a non-human instrument of detection and a concept is, at best, unknown. In so far as I am making the analogy between observation and detection, it would be controversial to suggest that instruments of detection have concepts or theories in the same way that humans (and other believers) do.[20] On the face of it, instruments of detection do not make judgements on the nature of what they detect. Non-epistemic detecting may be analogous to non-epistemic seeing, but given that the having of a concept is (might be) epistemically unique to humans, the analogy looks as though it might err in an important way. Thus, it is defensible that human detection has a special priority over other kinds of detection.

One solution to the problem would be to take an instrumental attitude towards our instruments of detection. Let me be clear on what I am not suggesting here. I am not suggesting that it is only *as if* the thermometer detects, *in the epistemic way,* what the measured temperature is. That it is *as if* the thermometer has some belief regarding what the temperature is, and that if the thermometer actually had that belief we would judge it as authoritative. Rather, the kind of instrumentalism I am offering is the kind where the thermometer – the instrument – *enables* us to measure temperature accurately. The previous chapter argued that the testing of an instrument was a genuine and necessary part of what it is to test more generally. This chapter has argued that we need not limit this kind of testing to just those observable features of the world. Our ability to test the accuracy of a thermometer does not depend on our ability to *observe* its accuracy. If this is right, the use of an instrument can extend our abilities.[21] There is nothing strange in claiming that having a concept required for 'observing' (construed in the broad scientific sense), might also involve the use of an instrument of detection.[22] There is a

[20] Controversial but not indefensible, see for instance Dennett [1981].
[21] Paul Teller [2001] develops a more restricted extension of what is observable. Teller argues that we are able to extend the range of observables using microscopes. It is not clear to me what might motivate such a restriction.
[22] The distinction being made here might be understood as the difference between Type II and Type III representational systems as outlined in Dretske [1988: ch. 3]. A Type II system derives its representational function from both human use (establishing a representational convention) and a causal relation to what is being represented. A Type III system is more autonomous than a Type II. For instance, Type III representational systems can misrepresent. A thermometer (Type II system) in so far as it *indicates* a temperature either succeeds or fails in this regard. Thermometers cannot misrepresent the temperature; they can only succeed or fail to indicate what the temperature is. Conversely, a type III system can misrepresent the temperature; maybe because of the failure of a thermometer to indicate what

way that van Fraassen might resist this extension of our abilities, and I will consider this defence in the next section; for the moment, let us consider one last type of objection.

It might be argued that how we test the relationship between empirical concepts that are used to infer just what has been observed might be different to the method of testing theories and concepts in my account of detection. Moreover, the theories used interpret instruments of detection are complex and move well beyond some fundamental understanding of the relation between empirical concepts. I do not see these issues as posing serious problems. Firstly, I am not suggesting that the method of testing theories required for detection is the final word, or that there are not alternative ways of testing. My claim amounts to this; we are able to test some detectable features of the world in the same way we might test what appears to be the case in the observable world. This may only get us to a certain point in our epistemic journey into the detectable world, but an important point none the less. Secondly, the theories by which we relate and understand empirical concepts are still relevant to complex theory construction and its relation to what is detected. Theories considering the detectable properties of the quantum world are strange precisely because of our conceptual understanding of the observable world.

Before we move on to consider how van Fraassen might respond, I will try to summarise this rather complex analogy. Observation, like detection, has both an epistemic and non-epistemic aspect. Of the epistemic part, both observation and detection depend upon the concepts involved and what theory you have when formulating judgements about what is being observed or detected. Of this epistemic part of observation, some beliefs regarding what is observable can be judged by testing theories that relate the relevant concepts used to form a belief, coupled with the non-epistemic aspect of what various modes of observing indicate. If what is indicated by different modes of perceiving is judged as being consistent (looks spherical, feels spherical, can manipulate as a sphere) we have reason to believe we have perceived correctly (a real sphere). In the case of detecting, we ought to believe our interpretation of what has been detected or measured in those contexts where, using a similar approach (the differential approach), we have tested what is indicated by our instruments of detection or operations of measurement. If what is indicated by different modes of detection or measurement is judged consistent, we ought to believe what is interpreted as being measured or detected.

3.4 Defending the Extension of Constructive Empiricism

According to van Fraassen's characterisation of the philosophy of science, on one hand we have the realists that push us towards a far right belief that science is aiming at true

---

the temperature is. Although Dretske's account is not without its problems, [Fodor 1990: ch. 3] this distinction is plausible enough for my limited use.

theories. On the other hand empiricists, because of 'epistemic modesty', would have us tending toward the far left: the belief that science only aims at providing theories that are true about what is being observed right now. 'Constructive empiricism finds an equilibrium point between the two extremes' [Monton and van Fraassen 2003: 407]. The equilibrium position is justified because the pay off for sticking your neck out and committing to the ontology of observables is that it provides a more credible account of the aim of science (empirical adequacy) for a small metaphysical price. On the other hand, committing to the full ontology of a theory is a high or unnecessary (ontological) cost, for the same (empirical) return. A constructive empiricist need only *accept* the full ontology of a theory while believing in the observables that allow it, the theory, to be empirically adequate. However, if what is detected pays an empirical dividend over what is observed, the ontological price of committing to a detectable world may in fact be worth the investment.

Although what it is to be observable remains largely unanalysed – this according to van Fraassen is a scientific issue – it is worthwhile comparing his position with Bertrand Russell's causal account of perception. According to Russell [1927: ch. 20] we can hypothesise that real objects cause what we perceive of them – their appearances. These appearances, caused by real material objects, Russell understood as 'groups of percepts' [Russell 1927: 211ff.]. Russell went on to suggest that we could synthesise some structural knowledge about the real object from a group of percepts. We will consider this proposal in chapter 6, but for the moment we need only appreciate that on Russell's construal, those things that, for van Fraassen, are observable, are, in fact, theoretical posits. This form of empiricism sees our epistemic position as something like Berkeley's idealism, where what we take to be real (by hypothesis) are the observable objects (rather than God) responsible for what appears to be the case. Russell argued that the causal account, when compared with solipsism, phenomenalism and Berkeley's idealism, was the best explanation for what we perceive. Constructive empiricists are not convinced that inferences to the best explanation are rationally compelling. So, given this scepticism, it looks like scientific theories can reveal what appears to be the case (how certain percepts are grouped) while we should remain agnostic about the explanation for such appearances (the observables). Empirical adequacy amounts to saving the phenomena, percepts, appearances or whatever, while we should remain agnostic about the 'observables' that account for the phenomena. How then might a constructive empiricist maintain the observable/unobservable distinction to avoid shifting towards a more conservative empiricism, while at the same time resisting my extension to what counts as empirical?[23]

We have already considered (in chapter 1) van Fraassen's response to a related difficulty noted by Grover Maxwell. The central problem Maxwell posed was this: how

[23] Marc Alspector-Kelly [2004] also draws some realist conclusions from van Fraassen's failure to resolve this issue.

might we draw the line between the observable, and unobservable? For instance, we can look through a pane of glass, spectacles, magnifying glass, microscope and electron microscope. At what point should we become agnostic towards the existence of an entity because it is unobservable? Van Fraassen claims he need not answer this question, for he can give clear examples of what is and what is not observable. It is true that the line is vague, but not much follows from this. Nor is it a response to say that we might evolve electron microscopic eyesight, thus enlarging the range of what is observable. We do not call a building portable just because we might evolve to be giants. 'Portable' is an ability word, and refers to us as human beings not giants, so it goes for 'observable'. Elsewhere van Fraassen has added that:

> laptop computers are portable and wine glasses fragile, even though some people are too weak to carry or break either. The limitations they refer to are of the human organism. So far, not even philosophers have suggested that the demarcation of the fragile has shifted after the development of such sophisticated instruments as the sledgehammer [van Fraassen 1992: 18-9].

'Detectable', he might add, is not about a human ability, or not just about a human ability. Extending what else counts as detectable has no bearing on what counts as humanly detectable (observable).[24]

We might note, however, that we do not have to wait until we evolve into giants to make buildings 'portable'. We can use equipment to make something portable. In fact, we do have portable buildings: caravans, mobile homes and portable offices on the penumbra of construction sites. They are portable because the technology used *enables* us as humans to move them about. Van Fraassen will of course claim that he is talking about a human ability, not the ability of a car to pull a caravan. Yet the point at issue is that a human ability is not always so easily detached from a technology. A person may be able to play a piece of music, but not without a musical instrument. Some people are able to hit golf balls long and straight, but not without a golf club. Closer to home: I have poor eyesight, but with my glasses on do I really have to maintain that I am still unable to see because of my reliance on a technology?[25]

---

[24] According to Musgrave [1985: 204ff.] a failure to see the relationship between what is observable and what counts as measurable or detectable, coupled with van Fraassen's concept of *acceptance,* results in him attributing to scientists a kind of schizophrenia. A constructive empiricist sees scientists accepting what they do not believe about the unobservable world. More pointedly, they do not believe what they have measured and detected of the unobservable world, even though they accept these results. It should be clear enough that my extension of constructive empiricism to include the detectable world would allow scientists to avoid this diagnosis.

[25] Musgrave [1985: 205] makes use of a similar point. We can consider the use of an electron microscope as extending our ability to observe things (regardless of what we might evolve into), and thus, what counts as observable.

Still, being observable might be like being fragile; it is a concept that does not depend on what technology is invented. There may be difficult or vague cases of what counts as an observation, and thus what is observable, but we can leave these cases for the epistemic community to resolve. Yet before we consider how the epistemic community might help here, we should have a closer look at the analogy between ‘fragile’ and observable. The first thing to note is that if someone were to point out that a certain set of wine glasses are fragile, what they are offering is a warning about how easy it is to break them. Once we understand that the real ability being referred to is breakability, then surely our range of breakable objects has been enlarged as a result of the development of the sledgehammer. My first job after graduating as an engineer was working in an underground mine as a grizzly-man: ‘grizzly’, presumably, because this unenviable job involved breaking rocks with a sledge hammer that were too large to be hoisted to surface. Rocks that were unbreakable were blasted with explosives at the end of the shift. In this context, what is breakable (by humans) has been enlarged by the invention of a sledgehammer. Perhaps van Fraassen could find another analogy to make the point. Yet if he is *just* talking about unaided human perception, why not just state this? The search for analogies would then be redundant. Unfortunately, van Fraassen cannot be this explicit.[26] To see why, we need to have a look at the role of the epistemic community in van Fraassen’s theory.

Van Fraassen gives the epistemic community two related goals. Firstly, the epistemic community has to work out the biology and physics of perception. This then would provide a solution to the second goal: to determine the ‘limitations to which ‘able’ in ‘observable’ refers – our limitation, *qua* human beings’ [van Fraassen 1980: 17]. But it is not clear how these two related goals might solve any of the difficult cases of what is in fact observable; at least not in a way that preserves his view on what is humanly observable. Consider a geologist looking through an eyeglass investigating the mineralogy of a rock sample. It is unclear whether the mineral structures observed using the eyeglass are observable (at least for a constructive empiricist). This seems like a case we wish our epistemic community to sort out. Yet it is not clear how the physics and biology of perception might help us with this problem. For a start, no science is required in order to sort out the limitations of the term ‘observable’, if what we mean by that term is ‘observable with an unaided human eye’. We need only look with an unaided eye to see what is observable. Unfortunately, demarcating the eyeglass situation as not relevant to unaided human perception hardly solves the problem as to what *is* observable in the human plus eyeglass case.

In recent work van Fraassen has argued that, when looking through an optical microscope, what is observable is a certain kind of image [van Fraassen 2008: ch. 4]. The argument for this is based on an analogy with what van Fraassen takes to be the

[26] In recent work, van Fraassen acknowledges that the line between what is observable and unobservable can be drawn in a ‘somewhat different way’ for different empiricists [van Fraassen 2008: 110].

relevant kinds of image – rainbows, mirages and what appears to be the bending of (straight) sticks immersed in water. On the basis of this analogy Van Fraassen concludes:

> The similarity I am pointing out to the rainbow (and which I propose as [a] basis for a possible way of thinking about microscopes) is not a similar lack of invariances that could be interpreted as invariances in *possible* relations to real objects ...

The 'invariances' here are the correlations we could make when observing rainbows, and presumably there are possible correlations to be made when observing through a microscope.

> The similarity I want to point to is just this: their [microscopes'] products are images; they are optically produced, publicly inspectable images. It is these images that are like the rainbow (they cannot themselves be represented as independent things). [van Fraassen 2008: 108]

Importantly, there are other kinds of images that do represent some real independent thing. Often images are images of something, the reflection of trees in water for instance. Rainbows, on the other hand, are 'produced', and thus do not represent any independent thing. Scientifically there is nothing wrong with a produced image; there is a fine tradition of experimentalism in the business of producing phenomena. Even so there is a distinction between this kind of produced image and the representative kind.

Although van Fraassen claims that we cannot see beyond appearances, he does outline how we might identify an image of the produced kind. This seems important given that rainbows can be reflected in water as much as trees can.

> Consider the rainbow. We realize pretty soon that there is no real material shining arch standing above the earth, although at first it looks that way. As a second guess we might think that certain parts of the clouds or haze are colored. But that cannot be maintained because if we move, we see the rainbow in a different location on that cloud or haze background. [van Fraassen 2008: 102]

Yet, presumably, a similar method can be applied to what is seen through a microscope. In fact there are clear cases where what we see through a microscope is an artifact (a produced image), maybe because of some fault in the lenses. But if we can identify this as an image produced by the faulty lens, how are we to contrast this artifact with the rest of the image 'produced' by the microscope? If, on the other hand, there is no distinction between the two images in the microscope case, what is it that makes the distinction when using the same method in the case of rainbows and trees?

What is relevant in the case of rainbows is the *appearance* of a shining arch located in the sky. I suppose there is no real problem in saying that a rainbow is an observable image, but this presupposes we understand the relationship between an appearance and observing that something is the case. Such an understanding would allow us to see why the appearance of the rainbow moving when we move had a bearing on the fact that the rainbow is a produced image. Appearances, then, are related to observing *that* something is the case, as much as the appearance of empirical adequacy is related to empirical adequacy. To this extent, one (theory-independent) fact that will need explaining by any empirically adequate theory is that, despite appearances, we can observe that there is no 'real material shining arch' moving in the sky as we move.

Returning to what counts as observable in our eyeglass case, the epistemic community can now investigate the relevant phenomena in light of what they judge to be the case. However, in order to determine just what is observable, some analysis of the instrument will be required. In particular, how the instrument functions in producing what appears to be the case. Maybe there are independent ways of checking what we interpret as indicated by this mode of detection (the differential approach). If there is, then the instrument *enables* the geologist to observe the observable – the mineral structure – using the eyeglass.[27] It is true that van Fraassen [esp. 1985: 297-9] has argued against a realist's interpretation of this kind of inference. However, if my analysis of observation is right, then simple perceptual judgements depend upon just this method.[28] If van Fraassen were to reject this account of perception, we are left wondering how we might make the ontological distinction between rainbows and trees, and thus how to assess empirical adequacy in light of what we observe.

Finally, a constructive empiricist might claim that, in the wash up, detection still depends on observation. All properties detected, in the end, need to be observed with the human senses in order to have epistemic surety. And yet we often believe what is detected over what is observed when using our senses alone. A thermometer provides a far more accurate guide to the temperature of a substance when compared with our unaided perception of it (e.g., our feeling how hot it is). We now have instruments that measure length, time, angle, mass and force to *imperceptible* degrees of accuracy.

[27] Hacking [1985: ch. 11] provides a good example of how the different modes of detection can arbitrate on this issue. Angus Menuge [1995] also makes use of our ability to 'cross-check' to enlarge the scope of what it is to observe.

[28] The inference that van Fraassen is arguing against here is something like an inference to the best explanation. I am not the first to notice that the use of this inference looks arbitrary if we can only use it when considering the observable world [Glymour 1985], or if rejected as a legitimate form of inference, may undermine his own account of empirical adequacy [Psillos 1999: ch. 9]. However, in light of the analogy I have made between observation and detection, the rejection of this form of reasoning looks particularly problematic.

3.5 Concluding Remarks

According to a detectionist, the term ‘observation’ means exactly what scientists usually mean by the use of this term: those detected or measured properties of world. If the above account is right, a reformed constructive empiricist armed with this usual sense of ‘observation’ can still maintain that if something is *observable* we are justified in believing it is real, and if it is not *observable* then we should remain a strict agnostic. Remember, a strict agnostic countenances as a real possibility that there is no unobservable entity beyond what is observed. Yet the failure to be detected (observed in the broad sense) provides a good reason, even for a realist on the far right, to remain agnostic about what is not (yet) detectable. It is true that some realists may still wish to consider the evidence as supporting, to some extent, the more hypothetical parts of a successful theory, and to this extent they hope to qualify this agnostic attitude. However, there remains a space of considerable agreement between what a reformed constructive empiricist might claim of the empirical (detectable) world and what a realist might claim. Rather than find an equilibrium point between empiricism and realism, a version of constructive empiricism can significantly reduce the distance between the two extremes.

Chapter 4: A Structure of Properties

4.1 Introduction

In the previous chapter I argued that detection had analogous features to observation and that, as a result, even a constructive empiricist could commit to the detectable world. In this chapter I hope to explore the kinds of thing we might claim as real on account of what we can detect or measure, according to the differential approach, described in the last chapter. To help explain what this form of realism amounts to I will introduce, and explicate, the concept of *a structure of properties*. It is this kind of thing that we can claim as *real* on the basis of information derived from what we detect or measure. A structure of properties is a collection of properties standing in certain physical, nomological or causal relations to each other. These properties might all be properties of the same entity, or they may be properties of different entities. Thus, as I construe matters, there are two senses in which we might use the term 'structure of properties'. Firstly, we can use this term to denote those properties of an entity that we consider to be in their important physical, nomological or causal relations to each other. Secondly, there are important physical, nomological or causal relationships between the properties of different entities. What we take as the relevant structure of properties depends upon epistemic considerations. Although I will not ignore arguments against how I wish to construe such claims, the central goal of the chapter is one of clarification. The ontology that I am developing here will be pertinent when we consider, in the following chapters, other forms of scientific realism that compete with mine. I should emphasise that, even if constructive empiricists were to reject how I wish to construe these claims, there remain good reasons for them to believe what science claims is detected of a detectable world.

4.2 Epistemic Realism

Given my emphasis on detection it should be clear that I am offering a version of what might be termed epistemic realism.[1] Epistemic claims about our ontology need not result in conflation or confusion. Indeed, as Hacking has noted, the whole project of scientific realism would be idle if we could not warrant what we thought existed as a result of our current best science [Hacking 1983: 28]. Following Newton-Smith's lead, Hacking provides the following 'ingredients' by which we might understand such an approach:[2]

(I) An *ontological* ingredient: scientific theories are either true or false, and that which a given theory is, is in virtue of how the world is.

[1] Chakravartty [1998, 2004] has also developed a form of structuralism, or 'semirealism' as he sometimes terms it, based on the relationship between detection and what is detected.
[2] These 'ingredients' were first presented in Newton-Smith [1978].

(II) A *causal* ingredient: if a theory is true, the theoretical terms of a theory denote theoretical entities which are causally responsible for the observable phenomena.

(III) An *epistemological* ingredient: we can have warranted belief in theories or entities (at least in principle) [Hacking 1983: 28].[3]

Hacking goes on to suggest that we can believe a particular entity exists without believing in the truth of any particular theory. This allows him to warrant a belief in the existence of an entity with little or no commitment to the truth of a theory: using an entity as a tool to intervene in the world will do. Thus, what he calls his 'theory realism' amounts to claims involving (I) and (III); while his 'entity realism' involves (III) and a slightly modified (II). Hacking modifies (II) because, apparently, 'one can believe in some entities without believing in any particular theory in which they are embedded' [Hacking 1983: 29].

I, too, think we can make some modifications to the above 'ingredients' such that we might make sense of claims regarding what has been detected or measured. I agree with Hacking in supposing that it is quite reasonable to claim that *entities* can be detected or measured. However, unless we consider what is detected of these entities, epistemic claims about them remain obscure. As we saw in chapter 1, according to Hacking's entity realism we can claim an entity exists in virtue of our interventions in the detectable world. Yet it is not clear, epistemically, what we are claiming of this entity in virtue of such interventions. To correct this, we will need to assess the descriptions made of the entities involved. If this is right we will have to attend to those theories that attempt to describe such entities.

We have already noted that the differential approach also allowed us to believe that our instruments or operations can measure or detect properties. It is a further question as to the nature of properties: for example, whether they are universals, or tropes or dispositions; and whether, and to what extent, causal powers, essences or natural necessities are involved. Even so, we noted that these properties can be found to be related when operationalising the theories that mention them. The fact that properties form these important relations or *structures* allows us to test one operational 'definition' against another. An epistemic realist can then justify claims about what is real based on *information* derived from the way the world is. Importantly, however, this information need not provide equal warrant for all parts of a theory. Van Fraassen was certainly right to note that scientists do believe that certain theoretical claims are true (or approximately so). But scientists may also wish to treat parts of their theories instrumentally, or may wish to remain agnostic about them, at least until they have more evidence.

[3] I have modified the numbering style.

In the next chapter we will consider several other forms of realism that attempt to partition theories into their important parts. To accommodate this possibility let us reconsider the above ingredients. The ability to assess *important parts* of a theory means that epistemic claims about what is real can have:

(i) An *ontological* ingredient: important parts of a scientific theory are either true or false in virtue of how the world is.

It is a central advantage of the differential approach to detection that it allows us to test important parts of a theory against each other. The test, we might presume, has a role in establishing what is responsible for what is detected or measured. If this is right we should also construe the term 'causally responsible' quite broadly so as to include detectable phenomena. Hence, epistemic claims regarding what is real can have:

(ii) A *causal* ingredient: if an important part of a theory is true, then those parts denote features which are causally responsible for the detectable phenomena.

As mentioned, for an epistemic realist, the warrant for beliefs about theories or entities or parts thereof is based on information about the world. Usually, what is warranted by this form of realism is informed by evidence taken from scientific practice. Actual examples of scientific investigation are meant to show how science itself provides the relevant information about what is in fact real. Anti-realists think no such information is available, or that the evidence counts against the realist. We will consider the anti-realist's evidence in the next chapter. Clearly, what counts as the 'relevant information' is importantly related to what we regard the information to be informing us of, and how it relates to important parts of our theories. These issues are a central concern of this chapter. For now we can adjust (III) to make explicit the role of information:

(iii) An *epistemological* ingredient: the relevant information can warrant belief in those important parts of theories or entities (at least in principle).

The type of epistemic realism I am developing involves all three ingredients. A *structure of properties* is the combination of properties that constitutes that part of an entity (or entities) that we have detected and/or measured, the results of which warrant the belief that an important part of a theory is true (or approximately so). Obviously such a structure (or structures) can involve complex causal relations with our instruments of measurement and detection. Apart from relating a structure of properties to important parts of our theories such that we can assess their truth, information, derived from a structure of properties, can warrant what a theory implies when applied in certain settings. That is to say, facts regarding a structure of properties can explain the security a theory enjoys. In what follows I will try and clarify these ideas with the help of some examples, as well as what others have said regarding related matters.

4.3 Empirical Structures

I wish to claim that in virtue of what we detect and measure we are provided with information about empirical structures or what I term a *structure of properties*. To get a better grasp of what these types of structures amount to, it will be useful to consider what van Fraassen claims constitutes an empirical structure. The point of doing so is that, by comparison with his account, I can explain what I mean by a structure of properties; it is not to argue against him.

After reviewing the history of various accounts of the role of structure in scientific theories, van Fraassen [2006] outlines what 'structure' means for an empiricist. The scene is set with the following claim, that according to *empiricist structuralism*:

> [T]here are (both in individual experience and in science) only two sorts of things we deal with directly. These are the concrete, observable things, events, and processes in nature on the one hand and on the other hand, the abstract structures studied in mathematics. We characterize the structure of the former in terms of the latter [van Fraassen 2006: 297].

This is a somewhat curious passage. It is not clear, at least by empirical standards, how we might deal directly with an 'abstract structure studied in mathematics'. Putting this issue to one side, van Fraassen is keen to emphasise that, apart from structure, a theory gets its '*credential*' from its empirical success, and that any transition from one theory to another has to involve an improvement in this '*credential*' [van Fraassen 2006: 302]. On the other hand, for an empiricist, structure is something that is preserved (and maybe even accumulated [van Fraassen 2006: 305]) through a change in theory. In so far as the old theories fit the data:

> There was something they got right: the structure, at some level of approximation, of those phenomena. Here the word 'structure' is used to point specifically to a certain character, defined by certain measurable parameters both old and new theory use to describe those empirical successes [van Fraassen 2006: 303-4].

Although it is vague what this 'certain character' is, we might note that if we don't maintain the relationship between the various measurable parameters we will lose some important information. For instance, we might measure various forces, masses and accelerations in an experiment. These measurements are perfectly legitimate as individual measurements, but, what is interesting is their relationship to each other. If, however, this relationship is not preserved then something is lost. It is presumably for this reason that van Fraassen talks of ‘low level laws’ or ‘simple equations’ that describe the structure of the phenomena.

Unfortunately, van Fraassen does not provide examples of what he counts as a simple equation or a low level law, but maybe what he has in mind is something like Einstein’s ‘simple equation’ used to calculate the angle of deflection of light passing the sun according to his *general theory of relativity* [Einstein 2004: 130]:[4]

1. $a = 1.7 \text{seconds of arc}/\Delta$.

What makes something ‘low level’ or ‘simple’ is not considered by van Fraassen. However, he does go on to provide a few suggestions as to what it is to embed the structural description of the phenomena (the low level laws) in the (more?) abstract structures of theoretical models. Still, it remains unclear what implications this is supposed to have. Does embeddability reveal our ability to summarise the data in various ways, or do the laws themselves refer to some real physical structure, distinct from the observable objects, that accounts for the organisation of the data? The answer is not clear.

The second, and related, issue involves van Fraassen’s understanding of the relationship between concrete observables, measurable parameters, data models, and phenomena. Van Fraassen uses some of these concepts interchangeably, but it is not at all clear how they relate to each other. Recently, van Fraassen [2008: Ch. 11] has developed further his account of empirical structuralism. Here he suggests that:

> [C]onstruction of a data model is precisely the selective relevant depiction of the phenomena *by the user of the theory* required for the possibility of representation of the phenomenon. [van Fraassen 2008: 252]

Pivotal in the above is what counts as ‘relevant’: and what counts as relevant, according to van Fraassen, is not determined by some abstract structure. Rather, the relevant depiction of the phenomena has something to do with the pragmatic considerations of the user of the theory. Here again, low level laws feature as summaries that fit the data or phenomenal experience [van Fraassen 2008: 247ff.].[5] However, in so far as laws fit what we select to be the relevant depiction of the data or phenomena, the question remains, is there some real structure that accounts for the way that the data and phenomena are organised according to those laws? As far as I can tell van Fraassen does not answer this question.

[4] In formula (1) ‘*a*’ is the angle deflection and ‘Δ’ is distance, measured in units of sun-radii, from the sun to a passing ray of light.

[5] These ‘simple equations’ and ‘low level laws’ van Fraassen now terms ‘surface models’ [van Fraassen 2008: 252].

For realists, our ability to select and summarise the data is a pragmatic consideration that does not, we hope, distort the role of a real structure responsible for the way the data is organised. Although the development of the data model is complex and interesting, typically, at some stage the data is organised into some sort of format with clear headings showing some relationship between each datum and the relevant calculation. Figure 1 provides a rather simple picture of what I have in mind.

Number of the Star.	First Co-ordinate.		Second Co-ordinate.	
	Observed.	Calculated.	Observed.	Calculated.
11 . .	−0.19	−0.22	+0.16	+0.02
5 . .	+0.29	+0.31	−0.46	−0.43
4 . .	+0.11	+0.10	+0.83	+0.74
3 . .	+0.20	+0.12	+1.00	+0.87
6 . .	+0.10	+0.04	+0.57	+0.40
10 . .	−0.08	+0.09	+0.35	+0.32
2 . .	+0.95	+0.85	−0.27	−0.09

Figure 1: A data model used to test the general theory of relativity. Taken from Einstein [2004: 132].

The calculated values in the above data model are determined according to equation (1). The equation itself is derived from Einstein's *general theory of relativity*.[6] Behind this data model, used to test Einstein's theory, is a larger body of data that includes angles, rotations, locations, distances, displacements, temperatures, times, periods, as well as corrections for all sorts of physical effects. All this data, accumulated at various times and places, had to be collated and synthesised into the above data model. Important here are the various theories used to locate all the equipment for the various experiments at the right time and place as well as ensuring that, as far as possible, the equipment functioned correctly. This is important knowledge to have when collating and interpreting a data model. In this process, an important part of the analysis will be the 'validation' or check of what we interpret as measured or detected, relative to both the organisation of the data, and the evaluation of each datum. As it turned out, a lot of data collected agreed with the effect due to *just* the Newtonian gravitational field. Importantly, these data were rejected. The details need not delay us, but the example does show that the construction of a data model is more than a collection of measurements.[7]

With this picture in mind, I claim, roughly, that each datum of a data model represents a determinate property. A structure that might be responsible for the composition of a data model I term *a structure of properties*. We have warrant for the belief that a structure of properties exists and is responsible for the composition of a data model if the data model is used to pass a differential test of just this structure.

[6] For more detail see Einstein [2004] or Eddington [1920: ch. 6]

[7] For some analysis see Mayo [1996b: 278ff.]; or for Eddington's own ruminations on the construction of the data model, Eddington [1920: ch. 7].

At this stage it is important to highlight some uses of the term ‘structure’ so as to avoid confusion about what is being claimed of a structure of properties. Sometimes we use the term ‘structure’ to pick out a certain token of a structure type; ‘that structure’ we say pointing. Sometimes, we use the term to refer to the type of structure involved. There is a difference between referring to a ‘concrete structure’ and referring to its structure.[8] To clarify the difference consider the following statement, ‘that structure$_{(token)}$, because of its unstable structure$_{(type)}$, is about to fall down’. A structure$_{(token)}$ can fall down, especially if it is unstable. On the other hand, it makes no sense to say that the instability identified as its structure$_{(type)}$ also falls down. Or, using another example, it is the structure$_{(type)}$ of a tennis ball – being hollow, elastic, spherical and pressurised with a gas – that allows a particular tennis ball – structure$_{(token)}$ – to behave in a certain way. The structure$_{(type)}$ these properties form, however, does not bounce with each tennis ball.

We can add to the picture the following: some properties do not bear a structural relation to others, or not the right type of relation. Of these we might say they form an accidental relation, or that they are merely correlated. All tennis balls might be a certain colour, but this property is not part of the structure$_{(type)}$ that accounts for the way it bounces, it is only correlated with that structure$_{(type)}$ of properties.

It should be clear that a structure need not be an entity or an object. A structure, in so far as it relates the properties of otherwise independent entities, can involve causal, physical or nomological relations. Such structures can have a structure$_{(type)}$ implicit in a law. On the other hand an entity or an object is not just structure. There is a form of structural realism that suggests that our best science (quantum based theories) may ask us to reconsider our commitment to entities in general, replacing them with a commitment to an all pervasive structure [Ladyman 1998, French and Ladyman 2003]. This is a controversial theory the criticism of which I leave to others [Psillos 2001, Cao 2003, Chakravartty 2003, Morganti 2004, Psillos 2006]. It should be clear that as I construe matters, a structure of properties does not replace an entity with just structure – no more than a bundle of properties might replace an entity with just a bundle.

It might seem that in so far as I am using the term *structure of properties* to pick out a ‘concrete’ structure it is merely operating as a synonym for the term ‘entity’. This suggestion misses the point of the term’s use. A structure of properties is what we can commit to as real, given the imperfect knowledge we have of what is real. The term ‘structure of properties’ refers to that part of the whole that we can pick out as real, based on the information we have. With perfect knowledge we might be able to replace all this structure talk with talk about unique individuals replete with their properties and relations to other individuals.

[8] Redhead [2001] also sees something like this distinction to be important when assessing structural realism.

Let me now pull together these ideas with some examples. Consider again some experimental measurement of the relationship between force, mass and acceleration (at reasonably low velocities). The experiment might involve a range of masses, forces and accelerations observed of various projectiles on an air track. Each relevant set of measurements involving a force, mass and acceleration is a representation of a structure$_{(token)}$ of properties that has the structure$_{(type)}$ implicit in the functional law $F=MA$. In this instance, the information regarding each force, mass and acceleration makes each relevant data set a representation of a structure$_{(token)}$ of properties, rather than an accidental correlation. It is in virtue of the information from testing for the causal relations between properties that we claim of a structure$_{(token)}$ of properties that it has a structure$_{(type)}$ implicit in the law $F=MA$.[9] The properties mentioned in a functional law are usually determinable properties, thus a functional law's truth-maker can be considered a structure$_{(token)}$ of determinable properties. Unless otherwise stated a structure of properties should be understood as a structure$_{(token)}$ of determinable properties. How individuals are related to this picture will be consider below (in §4.6), but roughly, individuals with dispositions are required to form structures of properties, though this need not fix, not without further theoretical consideration, any one of the determinable properties mentioned in the law. For instance, when considering the law $F=MA$, it might be supposed that the relevant individuals can decrease their mass and thus accelerate at a greater rate when a certain force is applied.

This distinction between individuals with dispositions and a structure of properties comes to the fore in contexts where we lack important knowledge about the individuals involved. Important in the discovery of the electron was J.J. Thomson's experimental work with cathode rays. We might ask, what did Thomson discover in 1897, when he calculated a mass to charge ratio for what constituted a cathode ray? It is true his evidence 'lends' itself to a particle view of cathode rays. Indeed his discovery of an electrostatic effect removed one of the obstacles to the particle theory. However, for the next two years, Thomson worked on developing methods for measuring the precise values for the mass and charge of the entities that made up the cathode ray. Although these determinations were met with some scepticism, ultimately he was proven right about the particle-like nature of an electron [Whittaker 1951: 364-5]. As more evidence was gathered, theorists were able to fill in more of the picture about the electron. It is clear that the results of Thomson's (and others') experimental work suggested that the particle view was the best explanation. But in 1897 whether the cathode ray was made up of one kind of individual with fixed mass and fixed charge, or a variety of kinds of individual with a fixed charge to mass ratio, or maybe a uniform substance, remained an open question. What then can we commit to as real given the information we have?

[9] Of course at high velocities the structure implied by such a law fails the test. Thus we can judge that the law, universally construed, is false.

In 1897 we had good evidence to suggest that cathode rays have a specific mass to charge ratio of 10^{-7} (grams/electromagnetic unit of charge).[10] As it happens, discovering the ratio involved measuring the deflection of the cathode ray under the influence of a magnetic and electrostatic field. In the experimental setting the following equation can be used.[11]

$$2.\quad m/e = \mathrm{H}^2\theta \mathrm{L}/\mathrm{E}\phi^2$$

Might we not believe that the part of the theory that implies that the constituents of a cathode ray have a mass to charge ratio (*m/e*) of 10^{-7} is true in virtue of the way the world is? Moreover, might we not believe that this structure of properties (the structure that ensures a fixed relationship between the mass and charge properties such that $m/e = 10^{-7}$) is causally responsible for what is detected or measured? This structure of properties is importantly related to the type of individual that constitutes a cathode ray. However, what is being proposed here also involves functional laws that relate this fixed mass-charge structure to those properties involved in detecting and measuring, for instance, those properties mentioned in (2). By considering the context of detection or measurement, and the functional laws involved, the range of properties involved is enlarged. I see no reason not to call this enlarged relationship between properties a larger structure of properties, and that such a structure is, in some sense, implicit in the relevant functional laws. This view need not result in any confusion between entities and the causal relations they enter into. If we are clear on what information is providing the warrant for the belief in the relevant parts of the theories or entities, we need not be misled by the term 'structure of properties'.

How then do we come by this important information? In the previous chapter, we saw how we might evaluate properties in virtue of our ability to detect, measure or observe. An important part of the evaluation involved having an independent check – the differential method. The check required two different theories that referred to the properties and two different instruments or operations over which the theories range. Whatever a measurable or detectable property is, in so far as is it is referred to in the theories used to measure or detect it, the implication is that it forms a structure with other properties mentioned in the theories. By relating these theories to each other we could check that what was indicated by each instrument or operation was consistent. If consistent, this justified the belief that we were detecting or measuring the same property. If not consistent, then something must have gone wrong. Thus we arrive at the rather common sense position that the numbers in a data base are evaluations of the

---

[10] The actual value is closer to 0.5686×10^{-7}. It is normally represented as a charge to mass ratio, and by convention, given a negative value.

[11] In the following equation m stands for mass, e for charge, H is the magnetic force, ϕ the angle of deflection due to the magnetic force, θ the angle of deflection due to the electrostatic force E, and L is the space over which the rays are influenced by the fields. This representation of the equation involved was taken form Harre [1981: 177] with some adjustment for clarity. The functional law can also be represented more generally as $m(d^2\mathbf{r}/dt^2) = e\mathbf{E} + e/c[\mathbf{v}.\mathbf{H}]$.

properties detected, measured or observed (in the broad sense). If differentially established, we can, with van Fraassen, describe this empirical structure with ‘low level laws’ and ‘simple equations’, but the method can establish more. It provides some ‘credential’ (or not) for the theories that, when combined to imply these low level laws, are tested by the data model. By assessing these empirical structures we either build up a theoretical description of the entities involved or test isolated parts of larger structures involving such entities.

I think what has been said so far clarifies the relationship between a structure of properties and the information we have in support of such a structure. In the next section I hope to analyse further what it is for information, regarding a structure of properties, to be implicit in a functional law. Pivotal here is the claim that functional laws provide information about what would happen under certain suppositions. I wish to defend the idea that this type of information is based on facts about the causal, physical or nomological structure of the world. To help clarify this idea I will consider an analysis of the laws of nature offered by Mark Lange. I think Lange’s analysis provides considerable insight into these matters. However, I hope to resist some of the commitments he thinks necessary for such an analysis.

4.4 Structure and the Laws of Nature

Functional laws (like (2) above) form, at least potentially, an important subset of the laws of nature. One influential analysis of the laws of nature argues that they depend upon natural necessities or a special nomological relation between property kinds [Dretske 1977, Tooley 1977, Armstrong 1983]. Typically the natural necessity involved is identified with a causal relation. Although I think it plausible that there is some form of natural necessity to causal relations, I also think, following Anscombe [1971], that the proposal requires some defence.[12] Nothing I want to claim here depends on such a defence. Nor will I consider a functional law’s relation to properties construed as universals.[13] Even so, I do suggest that functional laws, in so far as they mention determinable properties, are modal claims.

[12] For an attempt at such a defence see Tooley [1987]. It has also been argued that a reductive analysis of the modal character of the laws of nature in terms of something more ontologically basic than causal relations is destined to fail [Carroll 1987].

[13] There are reasons to believe that functional laws don’t involve universals. Insofar as a functional law mentions properties that might be considered as universals, then any determination of a determinable property is permitted. However, if any determination of a determinable property mentioned in a law is permitted, we will have to falsify or ‘truncate’ any other law that limits such a universal application. One solution to this problem is to consider, counter to the fact, one of the conflicting laws as false so as to judge what would happen according to the other [Armstrong 1983: §8.5]. Unfortunately there remain unwelcome consequences [Armstrong 1983: 171]. Another option might be to bite the bullet and, with Nancy Cartwright [1983], find the combination of functional laws as a local affair, grounded in its relation to a phenomenological law. According to Cartwright, any attempt to idealise functional laws into something universal may bear little relation to the facts.

Construing functional laws as modal claims provides us with information about what would be the case under different circumstances. Marc Lange has provided an important way of assessing this modal aspect of the laws of nature without presupposing natural necessities and possibilities. This provides considerable insight into how I wish to construe matters. However, Lange also argues that in order to explain these features of the laws of nature in a principled way we need to posit 'bare counterfactuals'. I will provide some reasons to resist this feature of Lange's, otherwise helpful, account.[14]

To get a grasp on the modal nature of a law, let us first consider Fred Dretske's [1977] account of the laws of nature. According to Dretske, empiricists traditionally believe that we can reduce a law of nature to a statement claiming a universal truth plus some other factor. Thus construed there is no real difference between a law of nature and any other kind of universal claim. For instance, an empiricist will represent the law that all diamonds have a refractive index of 2.419, as being equivalent to;

2. $\forall x\ (Dx \supset Rx)$

We should read (*2*) as: for all *x*, if *x* is a diamond (*D*), then *x* has a refractive index of 2.419 (*R*). Unfortunately this doesn't quite capture what it is to be a law. There seems some important connection between diamond-ness and refractive-index-ness of 2.419 missing. To see this, Dretske asks us to consider the universal statement; all dogs born at sea are cocker spaniels. If this were true, we could represent it as a universal truth with the same form as (*2*). Yet surely we want to claim that there is something about a diamond that *ensures* that it has a refractive index of 2.419, whereas, it is only accidentally true that all dogs born at sea are cocker spaniels. The difference is usually explained by claiming that laws are counterfactual 'supporting', whereas accidental yet true universal statements are not. That is, although it seems acceptable to say that, if this *x* were a diamond it would have a refractive index of 2.419, it is not acceptable to say that if this *x* were a dog born at sea it would have been a cocker spaniel. Dretske argues that laws support counterfactuals because they identify special property relations between property kinds. Dretske represents the relation as something like:

3. $D\text{-}ness \rightarrow R\text{-}ness$

I think much of this analysis is correct. However, Dretske is right to use scare quotes when arguing that laws of nature 'support' counterfactuals. Counterfactuals, it seems, can be obscure. Consider the following:

[14] In what follows I have chosen to ignore probabilistic laws, primarily because if a case cannot be made for a modal construal of simple functional laws, then there is no need to generalise such an account. But also, I don't think probabilistic laws would pose a serious problem.

> (A) Had the electron's charge been 5% greater, then the energy levels of the electrons within a silicon atom would have been 8% closer together, and so an electron would have needed to acquire less energy to ascend to a higher-energy orbital [Lange 2005a: 585].

Are there laws of nature that might support (A)? It's hard to say. The counterfactual seems to contradict a law (the one proposing what an electron's charge is). How then might we assess what it is that supports (A)? One response might be to claim that, given that (A) is a counter-legal claim, it would be too much to expect the laws of nature to support something that contradicts what is implicit in just these laws. Even so, (A) does seem to be making a claim relevant to the laws of nature. How are we supposed to view the relationship between the laws and this counterfactual?

Marc Lange has provided an analysis of the support between a law and its relevant counterfactuals within a more general account of counterfactual reasoning [Lange 1992, Lange 2000, Lange 2004, Lange 2005a, Lange 2005b]. Lange thinks that if we consider only those counterfactuals that are physically possible just according to the laws of nature, then the support that some natural necessity is supposed to provide between the laws and the relevant counterfactuals seems 'tailor made'. To see the problem, consider the following principle:

> *p* is physically necessary (i.e., a logical consequence of the laws of nature)
> if and only if
> $q \,\square\!\!\rightarrow p$ (i.e., "Had *q* been the case, then *p* would [still] have been the case") for every *q* that is physically possible (i.e., logically consistent with every logical consequence of the laws). [Lange 2005b: 417]

Lange notes that one problem with this principle is that 'it does not really show us a respect in which the physical necessities have an especially intimate relation to counterfactuals' [Lange 2005b: 417]. For a start, the principle uses modal terms that appear on both sides of the biconditional. Lange suggests that without some independent reason we lack an explanation for why we ought to 'privilege this particular set of counterfactual suppositions' [Lange 2005b: 418].[15] As we have seen above, some counterfactuals (such as (A)) may bear some important relation to the laws of nature, and yet they are excluded from the assessment by the above principle.

Lange also notes that the above principle 'is insufficient to determine whether a given truth is a law or an accident simply from its range of invariance under counterfactual suppositions' [Lange 2005: 418]. Some accidental truths are invariant under some counterfactual suppositions. For instance, the apparently accidental regularity 'all of the

[15] John Carroll [1987: 273-4] makes a similar point regarding the possible world account of laws.

apples on my tree are ripe' would have held 'had I worn a green shirt this morning or even had there been another apple on my tree. (Of course the accidental regularity might not still have held, had there been fewer sunny days last summer)' [Lange 2005b: 416]. We need some way of assessing the relationship between laws, accidents and counterfactuals such that we obtain the right result, but not just by stipulation. How might we avoid the threat of circularity?

Before we consider Lange's solution, there does seem a rather obvious way we might justify privileging the counterfactuals a law supports. We might suppose that we are attempting to describe, by the use of such laws, causal relations or causal powers of certain kinds of entities. Essentialists, for instance, will consider what is physically possible and what is of physical (or metaphysical) necessity in virtue of the kinds of thing that inhabit the world and the causal powers they have. However, Lange argues that this sort of approach is *ad hoc* when it comes down to the detail of choosing which counterfactuals are supported by such an approach [Lange 2004, Lange 5005a].

Lange's solution is methodological and involves assessing the *stability* of a set that may contain the laws of nature. The rough idea of *stability* is this:

> Let us allow … [any] set to pick out the range of counterfactual suppositions convenient to it, namely, the suppositions that are logically consistent with every member of that set. If, under each of these suppositions, all of the set's members would still have been true then the set qualifies as stable. [Lange 2005b: 419]

If a set contains laws and only laws it will be stable; that is, its truth is invariant under any supposition consistent with it. If the set contains an accidental generalisation then the set will be unstable – variant under suppositions consistent with each member of the set. Consider the set that contains the fact 'all of the apples on my tree are ripe'. We might first attempt to test the set for stability by considering the supposition 'had I worn a green shirt today' in relation to 'all of the apples on my tree are ripe'. If the counterfactual formed is true, this test reveals that the set might be stable. Now consider the supposition 'had there been less sunny days last summer'. Although consistent with the fact contained in the set under question, when this supposition is combined with the fact, the counterfactual formed may not be true. Hence the set is not stable.[16] Suppositions that contradict a generalisation contained in a set can still be assessed. However, such an assessment has no bearing on the stability of the set. Finally, one must be careful not to import claims that might rig the result. So the claim 'all electrons are negatively charged' is allowed as a member of a potentially stable set of laws, but

[16] Lange gives further details to show that his methodology gets the right result [Lange 2000, Lange 2005b]. Moreover, the results are quite general as the stable set approach also explains grades of necessity.

not ‘it is a law that all electrons are negatively charged’ or ‘it is an accident that all electrons are negatively charged’.

One of the advantages of this approach is that stability does not give any special role to the counterfactuals ‘supported’ by the natural laws. Thus Lange’s concept of stability can avoid the accusation of circularity. According to Lange we have a way of determining what it means for a truth to possess the relevant kind of necessity, without any apparent appeal to just that kind of necessity. Yet, even though Lange’s methodology might get the right result, we are still left wondering what it is that enables us to form a stable set. Clearly assessing stability requires the ability to assess counterfactuals. In some places Lange claims neutrality: laws might support counterfactuals or maybe the other way round – ‘perhaps counterfactuals (relativised to a context) are as ontologically basic as actuals’ [Lange 2005b: 247]. Perhaps ‘something else is ontologically prior to laws and counterfactuals, serving as a truth-maker of each, and thereby putting laws and counterfactuals into the relation I have described’ [Lange 2005b: 427]. Elsewhere, Lange wants to make a stronger claim:

> These counterfactuals are ontologically prior to the facts about what the laws are. Whether this direction of ontological priority encounters greater problems than it avoids is a question for further research to address [Lange 2004: 240].[17]

Before we assess this claim we might see if Lange’s approach can accommodate the way in which I wish to construe a functional law. I wish to construe a functional law as a modal claim precisely because of the counterfactuals it implies. Consider the law $F = MA$. We can enumerate this law to form a counterfactual. ‘Enumerated’ in this context means providing the appropriate values for the theoretical terms mentioned in the functional law. For instance, according to the law $F=MA$, if a force of 2000 Newtons were to be applied to a 1 kilogram (unrestrained) mass it would accelerate at 2 metres per second (for as long as the force is applied). All similar counterfactuals formed using the functional law are just mathematical enumerations of the law itself.

Lange’s approach adds considerable insight into the modal nature of a functional law. In order for a functional law to be a member of stable set we need to strip it of its law-like implications. Hence, let us consider a functional ‘law’ as just a convenient way of summarising the facts. With van Fraassen we can call such an interpretation ‘a simple equation’. As a summary of some facts, the equation does not have any implications regarding other counterfactual ‘evaluations’ consistent with it. Thus understood, a simple equation that might be a functional law is ripe for counterfactual assessment in a (potentially) stable set of laws. If the set is stable then we will have found those counterfactuals consistent under each interpretation of the simple equation (and other

[17] See also Lange [2005a: 587]

laws) as true.[18] As Lange mentioned, the stable set approach is quite consistent with there being something ontologically prior to laws and counterfactuals, acting as a truth maker for both and thus explaining their relation to each other. To this extent I claim a structure of properties is a defensible contender for this role.

All this seems helpful. What then of the issue of ontological priority of counterfactual truths suggested above? Although interesting, I think this issue of primitive counterfactual truth is largely independent of the central insight Lange has provided for my account. Even so, I think there are reasons to avoid primitive counterfactuals. Lange himself admits that they are somewhat strange [Lange 2005a: 587]. There are, however, reasons to suppose there are such truths. As we saw with (A) above, scientists sometimes do engage in counterlegal reasoning that conflicts with the laws of nature; maybe we do require primitive counterfactuals as a basis for this sort of reasoning. However, the main reason for postulating bare counterfactuals, it seems, is that with them we can explain the basis for modal reasoning in a principled way. As mentioned above, Lange thinks the approach offered by the essentialists is *ad hoc*. This is because essentialism just maps modal reasoning onto whatever science discovers about the essential properties of the natural kinds that inhabit, or could inhabit, our world. ‘Essentialism burdened by this sort of ad hoc fine tuning has merely had the right answer inserted into it by hand’ [Lange 2005a: 583].

Yet it is not clear to me why it is principled, and thus not *ad hoc*, to propose that ‘counterfactuals lie at the bottom of the world and are partly responsible for the laws’ [Lange 2005a: 587]. Nor does it seem to me that the semantics of counterfactual reasoning support there being primitive counterfactuals, the use of which is principled and thus not *ad hoc*. Let us consider (A) again. We might begin by noticing that, even with a stable set of laws, we cannot assess the truth of (A). So it looks as though (A) may depend upon bare counterfactual truths. But before we grant this, let us suppose that some scientists think that (A) is false, and that they considered the following true:

> (A′) Had the electron’s charge been 5% greater, then the energy levels of the electrons within a silicon atom would have been 16% (not 8%) closer together, and so an electron would have needed to acquire (even) less energy to ascend to a higher-energy orbital.

Now, in order to settle these matters, would it not be considered *ad hoc* if one group of scientists appealed to bare counterfactual truths in support of their position? Might they not be accused of having inserted the right answer into their thesis by hand? I suggest that if the issue could be sorted out, then it would be done so by appealing to the relevant laws of nature, background assumptions, theoretical commitments and the

[18] It might be the case that if we considered the complete set of functional laws the counterfactual supposition consistent with all of them may be quite limited.

facts. There would be counterfactual reasoning, no doubt questioning these background assumption an theoretical commitments, but we might presume that they would at least try, as far as possible, to stay within the implications of the relevant laws and facts. Of course there is at least one ‘primitive’ counterfactual, the one regarding the different charge we might presume an electron to have. But although, for the sake of reasoning, this counterfactual is supposed to be true, it is also believed, *in fact*, to be false. And it is this fact that makes it difficult to assess the basis for such counterfactual reasoning. The important point of this example, then, is not to settle what constitutes the basis for such counterfactual reasoning. The point of the example is to question why, when reasoning under such suppositions, an appeal to primitive counterfactuals would be considered principled over what others might take to be primitive.

Counterlegal reasoning seems to me to be rather exceptional in scientific practice. In the more usual cases of scientific reasoning there is no doubt which way the support runs. Consider the design context where the treatment of functional laws as modal claims is common. Whether designing an experiment or designing an artefact, the designer can see, in virtue of the enumeration of the relevant laws, what effect various enumerations of the theoretical terms in a law would have on the design. Such evaluations may motivate adjustments to the design, maybe so as to avoid a particular problem the application of the laws envisage (a calculated stress exceeding a material strength). Importantly, many designs are usually produced with only a small subset of them, if any, realised. It is not unreasonable to consider these unrealised designs as counterfactuals. Yet we don’t consider the unrealised designs, even if faultless, as supporting, in any way, the laws used in their construction. The assumption of this reasoning seems to be something like; a law being true makes it the case that certain counterfactuals are true. I am not suggesting these semantic issues settle the matter as to which way the support runs, and thus what it is that is more fundamental. Even so, any account of primitive counterfactuals supporting a law will need to accommodate the direction of support that seems implicit in this form of reasoning.

What about counterfactual reasoning involving accidental regularities, do we need to appeal to primitive counterfactuals as a basis for this reasoning? Consider again Lange’s proposal that it is an accident that ‘all of the apples on my tree are ripe’. How accidental, we might ask, is this regularity? Let us say we hold some (vague) theory about what causes fruit to ripen and how this relates to the seasons of the year. As a result we may believe that certain relevant causal relations underpin the regularity. This then seems a plausible basis for the invariance (and variance) we find in the accidental regularity under various suppositions. For instance, the relevant causal relation provides a basis for the invariance of the accidental regularity under the supposition ‘had there been another apple on my tree’, and its variance under ‘had there been fewer sunny days last summer’. Importantly, because our assessment uses the stable set approach we understand that, although law-like to some extent, the regularity is not a strict or fundamental law, and thus, and to this extent, accidental. Accidental regularities, it seems, come in degrees of law-likeness. Consider another accidental correlation, 'all tosses of this coin turn up heads'. This accident, we might suppose, is less stable than the accidental regularity 'all of the apples on my tree are ripe'. We can justify this belief by noticing that it fails an important and analogous test to the one used to test the accidental regularity 'all of the apples on my tree are ripe'. We might not believe that, 'had I tossed the coin one more time' then it would be the case that 'all tosses of this coin turn up heads'. What might count as the basis for such an assessment? Might we not claim the reason it fails the test has something to do with the relevant causal relations underpinning this accidental regularity. Notice that in both cases we used the stable set approach to assess the accidental regularities; indeed, it provided an explanation for why we might believe one accidental regularity as more accidental than the other. Yet we did not need to appeal to primitive counterfactuals in this assessment. Lange might suggest that the introduction of 'relevant causes' in the above account is *ad hoc*. Moreover, what counts as a relevant cause is itself ripe for counterfactual assessment, and thus its introduction merely delays the issue.[19] Fair enough, I suppose. But what if, rather than appealing to the relevant causes, appeals were made to primitive counterfactuals? It is not obvious that this would make the reasoning any more principled.

The stable set approach is doing the real work of providing a principled approach when assessing the relationship between laws and counterfactuals. This, I think, is the central contribution Lange has made. The need for primitive counterfactuals that might make the approach principled in a way that essentialism (or some other account) is less convincing. Of course Lange may have independent arguments to support the need for such primitive counterfactuals. As it stands the stable set approach does not imply the need for these particular primitives.

It seems that the stable set approach helps in the explanation of the relationship between functional laws and their relevant counterfactuals. This leaves it open for a structure of properties to be a basis for functional laws, and the counterfactuals they imply. We can of course discover significant limitations regarding the truth of what a functional law implies. This assessment will be judged in relation to the accuracy of the relevant measurement or reliability of the instrument of detection. These assessments will depend on a range of secure, and not so secure, assumptions, background theories, intuitive concepts, and the relevant facts. To this extent, facts regarding these structures, (coupled with the relevant background theories, assumptions and concepts) provide warrant (or not) for the functional law, and thus what the law implies.[20]

[19] Psillos (2004) argues that, although it is not without its problems, the counterfactual analysis of causation is 'more basic' when compared to its mechanistic rivals. Even so, Psillos argues that, in practice, it is a combined approach that offers us a better understanding of the nature of causation.

[20] It might seem to some that this kind of appeal to a scientific practice to settle these matters trivialises the philosophical analysis of the laws of nature. I don't think the criticism sticks. For a start, I don't think scientific practice is philosophically neutral. Nor am I alone in suggesting science might be involved in

4.5 Information Regarding a Structure of Properties

Lange's approach has provided considerable insight into the informative nature of functional laws. By analysing those suppositions consistent with a functional law we could see, in virtue of the stable set approach, just how informative these laws are. Functional laws, especially, inform us about an enormous range of counterfactual circumstances. This is particularly important when designing on the basis of such laws. The information implicit in these laws, I suggest, is based on important facts regarding certain causal, physical or nomological structures of properties. All this analysis would be undermined, however, if these important facts turned out, upon analysis, to be facts about the way we impose concepts or theories on certain regularities. The nature of causation, as well as what counts as physical or nomological relations, is a highly contested philosophical terrain. Although what I am proposing is meant to be somewhat neutral in regard to the detail of these accounts, my position is obviously aligned with the realist's interpretation of such structures.

In this section I will consider what it is to be epistemic realist about those structures that constitute causal relations. I have argued that, if we wish to make an epistemic claim regarding the existence of a causal relation, we need some account of the information upon which we might base such a belief. However, some accounts of causation consider that there is no objective grounding to causal relations, or that we have no reason to believe that there is a real basis to such relations.[21] According to these views, what counts as a causal relation amounts to certain conventional practices, psychological attitudes, habits or other pragmatic interests we have in certain regularities. As a realist I can certainly recognise the potential for these aspects to feature in the causal relations we identify. However, I maintain that there is some information regarding a causal relation to be had independent of these factors. Given these broad views, it might seem that there are two related issues: what information is there to be had regarding a causal relation, and secondly, how to represent such information. The problem for the epistemic realist is that, when we consider a law supposedly based on a causal relation, the information provided seems quite compatible with a sparse view of such relations.

To clarify the problem, we might consider the criticism van Fraassen [1987] makes of Armstrong's representation of necessary or nomic relations between properties.[22] Van

---

sorting out some of these modal issues; see for instance Menzies [1993], Ellis [1996] or Clendinnen [1999]. Nor should I be interpreted as trying to simplify a complex area.

[21] Max Kistler [2006] attributes this view to Hume. 'David Hume has raised a powerful challenge to the idea that causation has an objective grounding. Any analysis of this concept that aims at giving an objective value to statements must answer to Hume's objections' [Kistler 2006: 11].

[22] Van Fraassen is one of several who line up as modern sceptics of natural necessity [Carroll 1987, Lewis 1983, Mellor 1980].

Fraassen asks us to consider the following schema that he thinks is representative of Armstrong's theory:

4. $N_1(F,G)(a$'s being F, a's being $G)$

We can read this as, a's being F and a's being G is ensured by the necessary relation N_1 between F and G. We can generalise this necessary relation as:

5. $N(F,G)$

However (*3*) and (*4*) say nothing about other particulars. Given that the particular chosen should not matter, the following seems reasonable:

6. If b is (F and G) then $N_2(F,G)(b$'s being F, b's being $G)$

In (*6*) any object that has F and G can be b. So we seem to have a way of capturing the necessary-relation for any instance of F and G of any particular that has these properties. But then nothing much follows from this. For what we really need to derive from (*6*) is the universal statement: all things that are F are G. What we have by virtue of (*6*) is only that, of the particulars that have F and G, the relation between the F and G will be a necessary one. The identification of a necessary relation is compatible with the assignment of necessity *after* the individuals that instantiate the F and G have been sorted. The problem seems to be that the N relation does not do any informative work outside of those particulars that already instantiate F and G. Even if we could identify such a relation, van Fraassen asks, 'what information does the statement that one property necessitates another give us about what happens and what things are like' [van Fraassen 1989: 96]?[23]

Armstrong, as we will soon see, proposes that we can identify singular causal relations that instantiate a relation in which F necessitates G. However, I am not defending a natural necessity account of causation. Even so, van Fraassen's point could be seen to count against what I claim of a (causal) structure of properties. If we interpret $N(F,G)$ as identifying a $\text{structure}_{(\text{type})}$ of properties then, upon analysis, it seems we need only construe this structure as the regular occurrence of two properties certain objects regularly instantiate together. So construed, we might now ask, what information over and above the regular co-occurrence of the properties F and G does $N(F,G)$ imply? As I do hope to claim that there is more information to be had regarding a structure of properties, van Fraassen's question seems pertinent.

[23] Matters are somewhat more complicated than what is suggested here. Van Fraassen [1989] identifies two problems 'the identification problem' and 'the inference problem', by their relationship to each other they create a dilemma for a 'universal' theory of laws. However, it seems it is the inference problem that generates a problem for Armstrong's theory [van Fraassen 1989: 103-4].

In response to van Fraassen's criticism, Armstrong proposes that we 'regularly do perceive that one thing causes another … A conspicuous case being the perception of pressure on our body' [Armstrong 1993: 421, see also Armstrong 1988]. This experience allows us to identify singular causes. If this is right, Armstrong asks '[m]ay we not hypothesise that this uniformity holds *because* something's being F *brings it about* that the same something becomes G' [Armstrong 1993: 422]? The higher order laws are thus atomic facts that relate *types* rather than *tokens* of causal relations, and if hypothesised as an explanation, imply all the (non-defeated) token causal relations. The relevant necessary relations are higher order explanations that, if accepted as an explanation, allow an analytic or conceptual inference between types of cause and types of effect. This allows a necessary relation between universals to both entail and explain regularities.

I think this response is fine as far as it goes, but before we consider another response, we might provide a bit more colour to this experience of causal relations that Armstrong has identified. Although feeling a pressure may be an instance of perceiving a causal relation there seems more to it than this. For instance, using a torque wrench to tighten a wheel-nut is a rich experience of causal relations that can involve feeling a pressure. Importantly, not only do we perceive the causal relations involved but we are also using them, in this particular instance, to *tighten* a wheel-nut. We fail or succeed in our attempt to tighten a wheel-nut according to our ability to identify what the relevant causal relations are, and how proficient we are in using them. If there were no such things as causal relations, no success or failure would be possible; though we could, I suppose, talk vacuously about our ability to tighten wheel-nuts. The context in which we use causal relations seems more of a challenge to the sceptic than the context of just our perception of them.

Van Fraassen's analysis suggests that when analysing laws involving *F*'s and *G*'s we find no relevant content bearing *N* relation, and thus we are left wondering about the informative nature of this relation. On the other hand if some causal relations do carry content, for instance '__ tightens __', we might ask, wherein lies the *N* relation? Presumably Armstrong would suggest that, if *a* tightens *b* is identified as the relational property *Fab*, and *Gb* the resulting fastened *b*, the natural necessity $N(F,G)$ can account for this and other similar singular causal relations. Armstrong thinks natural necessities operate at a higher level than singular causal relations and thus need not carry the content of the singular causes. Still, there is some mystery as to the role of the *content* of a causal relation when accounting for its effect.

This mystery may only be a feature of higher order laws involving F's and G's. Functional laws seem quite adept at accounting for the content of causal relations between properties. In so far as a functional law mentions various determinable properties it can inform us about the causal, physical and nomological relations these properties enter into. Consider some of the more usual examples of the functional laws of nature.

7. $n_2/n_1 = \sin\,\theta_1/\theta_2$ *(Snell's Law)*
8. $Uf - Ui = Q - W$ *(First Law of Thermodynamics)*
9. $F = q_1q_2/4\pi\varepsilon_o r^2$ *(Coulomb's Law)*

If we allow these functional laws to be construed modally then they can *inform* us about would happen under various evaluations of the properties involved. Not only do they inform us about what will happen given current and future circumstances, but what would happen in circumstances that may never be realised. As we have seen, the design setting reveals the kind of information available from a functional law. They, the laws, inform us of what would happen in the circumstance envisaged by the design. Insofar as these laws are about a causal, physical or nomological structure their modal nature carries a lot of information about these important relations.

Contrary to what I claim, the regularity view of a functional law understands it as a sophisticated, although only convenient, summary of a correlation in the data [Forge 1986, van Fraassen 1989, Smart 1993,].[24] Regardless of how well a functional law fits the data, the information required for designing is not just about a correlation in some data. This is because it is not clear that enough information, for the purpose of a design, is available from a summary of the data. In the design setting we are not so much interested in 'what happens and what things are like' as much as what would happen given different circumstances – the circumstances the design envisages. Even if it is claimed that in the experimental context, experimenting is theory construction by other means [van Fraassen 1980: 77], this is not the kind of theorising that is done in the design context (which will be a part of the experimental context). In the design context certain functional laws are used to design with, not theorise about. Without information about the relationship between determinable properties mentioned in the laws, designing would be difficult if not impossible. Nor is this information about conventions and psychology. The correct use of a functional law in a design can avert disaster, and not just a linguistic or logical disaster. For instance, the use of functional laws might reveal a problem with the load distribution in an aeroplane wing under certain flight conditions. Accordingly, we can adjust the design to avoid the problem. The use of functional laws can inform us of, not only what happens and what things are like in certain circumstances, but also what would happen, even if these circumstances (thankfully) may never occur.

Pointing to the modal interpretation of functional laws in the design setting certainly does not settle the matter as to how we are able to obtain information about a causal structure implied by a law. I think that the differential approach could form an important

[24] Van Fraassen [1989] argues there are only regularities and thus no laws worthy of the name.

part of the story, but a more detailed account of causation will also be required. Pointing to the use of laws in the design setting does, however, make it clear the kind of information available, and how it might be represented in a functional law. Furthermore, it is unlikely that the information derivable from a functional law is based on facts about conventions imposed on regularities or convenient ways of summarising data. Given that we can evaluate functional laws in a way that informs us of what would happen in counterfactual circumstances, we can see how such content, regarding a structure$_{(type)}$ of properties, might be implicit in such a law. There seems no reason, then, why a realist can't have information regarding what they take to be the causal structure$_{(type)}$ of properties implicit in a functional law.

4.6 Individuals and Structures of Properties

So far I have focused on information based on facts regarding a structure of properties and how this is related to what is implicit in functional laws. Yet we do want to make a distinction between the structure$_{(type)}$ of properties implicit in functional laws, and that part of an entity responsible for what we detect or measure. We do wish to say that the entities that constitute cathode rays have a specific mass to charge ratio that is responsible for what we detect and measure. However, in order to detect and measure what constitutes a cathode ray, we need to consider the properties of a cathode ray in relation to a larger structure of properties. Fortunately there is a straightforward account that allows us to articulate this distinction. Although sceptical of certain metaphysical posits – including possible worlds, universals, natural necessities and essential properties – Hugh Mellor [1980] has developed an account of individuals with dispositions that can act as truth makers for the laws of nature.[25] According to Mellor, in a Newtonian world we can claim of an individual *i* with a mass *m*:

(B) If *i* were subject to a force of *f* newtons, it would then accelerate in the direction of the force at *f*/*m* metres/second^2 [Mellor 1980: 120].[26]

What is important for the truth of this conditional is that there has to be some individual *i* with a mass *m*. It is in virtue of the dispositions of *i* with a mass *m* that it can be forced and accelerated according to the law. One of the virtues claimed of this approach is that it avoids problems associated with vacuous laws. If laws are distinct from the individuals that instantiate them, this seems to imply that there may be laws that govern interactions between individuals that do not exist. These uninstantiated laws make it difficult to assess what it is that makes the law true or false.[27]

[25] See also Mellor [2000].
[26] This conditional is slightly modified for clarity.
[27] More detail on the nature of the vacuity is provided in Carroll [1987], Lewis [1983], Mellor [1980] and van Fraassen [1987]. Armstrong [1983: 114-5] offers some tentative solutions to the kind of problem highlighted, though acknowledges they are not entirely satisfactory.

To see the problem we might consider an example offered by John Carroll [1987]. We might consider two universes: $universe_1$ and $universe_2$. In $universe_1$ there are X-particles that interact with W-fields according to a law L_1.[28] There are however, no Y-particles or Z-fields in $universe_1$. In $universe_2$ the situation is reversed, with law L_2 governing the interaction of Y-particles in Z-fields. In $universe_1$, an X-particle has the property spin-up when interacting in W-fields according to L_1. In $universe_2$ spin-down is ensured of a Y-particle interacting with a Z-field according to L_2. The central problem is that 'the Nomic Realists cannot say in an illuminating way what makes it the case [that] the appropriate relation obtains in the first universe, but not the second' [Carroll 1987: 265] (or vice versa). There is nothing in the accounts of the two universes that *explains* why the different laws are instantiated. Appealing to a necessary relation between the properties X-particle-ness, W-field-ness and spin-up-ness only delays the problem: what is the nature of this necessary relation and how does it help? Is there any further illumination if, rather than claiming merely that the particles *do* obey their respective laws in each particular universe, we claim further that the particles do obey their respective laws *necessarily*.

Mellor's account avoids these problems as he need not consider what truth bearing relation there is between laws and non-existent X-particles and the like. Mellor does not need to consider how laws apply in such vacuous cases as he does not believe that laws hold in virtue of natural necessities that relate universals.[29] According to Mellor's account, the laws are made true or false by individuals in this world. Interestingly, Mellor does note that once we have all the various reduction sentences (such as (B) above) involving functional laws ranging over the theoretical term mass, we can 'distinguish all masses from each other and from other factual properties' [Mellor 2000: 770]. This seems a bold claim that, if true, would presumably *determine* what would happen according to the laws pertaining to each '*i* with a mass *m*'. For some, to account for this determinism the individuals involved might be thought of as having more than just dispositions: causal powers or essential properties might fill this role [Shoemaker 1980, Ellis & Lierse 1994, Mumford 1998: ch. 10, Ellis 2001]. The details here need not delay us. We need only note that having individuals with dispositions (and whatever else is required) allows us to garner information about the empirical structure of the world. We can reason counterfactually on the basis of this information, while leaving the individuals with their dispositions as truth makers for the functional laws (and other theoretical descriptions). The kind of reductive analysis that Mellor has provided shows us that the relationship between functional laws and the individuals involved need not result in any confusion regarding epistemic claims about what is real.

How might individuals with dispositions feature in an epistemic realist's interpretation of a theory when we don't have a complete picture of the individual entity involved?

[28] I have slightly modified the example for clarity.
[29] Nomic realists have responses to such cases [Armstrong 1983: ch. 8].

Given what we currently detect or measure, it seems reasonable to consider a structure of properties as a 'shell' or 'frame' for the actual individual(s) involved. We can of course claim that some unique individuals with their dispositions must exist such that we can account for what we have detected and measured. But confining epistemological claims regarding what is real to descriptions that might uniquely identify those individuals that make up the kinds of entity of the world seems too restrictive. As Chakravartty notes:

> [H]owever realists choose to construct particulars out of properties, they do so on the basis of a belief in the existence of those properties. That is the bedrock of realism. Properties lend themselves to different forms of packaging, but as a feature of scientific description, this does not by itself compromise realism with respect to the relevant packages [Chakravartty 2003: 874].

The point being made, we might presume, is that we can discover how properties are differently related, but this does not undermine the fact that they have to be related in some relevant package. This then allows us the belief (maybe because of our imperfect knowledge) that different kinds of individual could have instantiated the relevant properties, but not the belief that no individuals were involved. This information, given what else we know, can be used to determine what would happen to different individuals that might instantiate a structure$_{(type)}$ implicit in a functional law.

Just after his 1897 experiments, as far as Thomson knew, different types of individual (equipped with different dispositions, causal powers and whatever) might have constituted a cathode ray so as to form a constant charge to mass ratio. On the other hand, the claim that properties 'lend themselves to different forms of packaging' somewhat understates matters. Thomson, in his 1897 experiments, discovered something about the 'packaging' of the properties that constitute a cathode ray. Certain properties, mass and charge, come packaged together in a fixed, discoverable relationship. Regardless of future discoveries, after 1897 we knew something about this structure of properties that constitutes a cathode ray and we knew what we would detect and measure in virtue of such a structure. The warrant for such claims is based on information regarding a larger structure$_{(type)}$ of properties that included the relevant properties of the instruments of detection and measurement. It is just this structure$_{(type)}$ of properties that is implicit in the functional law (2) above.

4.7 Concluding Remarks

When we theorise about the individuals under scientific investigation we might use a range of analogies, familiar concepts, diagrams, graphs, the methodology of paradigm experiments, functional laws, background assumptions, and so on. All this provides a rich picture of the entity under investigation and no doubt guides the development of the theory itself. As epistemic realists, I suggest we can compartmentalise this picture and claim which parts correspond to some real feature of the world and which parts we should remain agnostic about. Unlike constructive empiricism, however, we need not remain agnostic for too long. The epistemic reach of our instruments of detection and measurement grows daily. And as long as there is information to be had, we can be epistemic realists about what we are informed of by the differential approach to detection. This chapter has attempted to provide something this information can be about. In the next chapter we will consider what other realists have made of the concept of a structure and its relation to certain scientific theories.

Chapter 5: Differential Realism

5.1 Introduction

The overall aim of this thesis is to defend a form of scientific realism that provides justification for a realist interpretation of important theoretical terms and the approximate truth of important parts of our best theories. When a scientist claims that DNA is now understood as a complex molecule responsible for cell development, a scientific realist (of the type I am defending) claims that we can be justified in believing that there are things in the world that more or less correspond to what we mean by the use of the terms 'cell' and 'DNA'. Moreover, that the relevant properties these entities have or instantiate, form causal, physical or nomological relations that are, in virtue of what we detect or measure, describable in truth-assessable forms. According to the anti-realist, however, the problem for a realistic account of science is that, upon analysis, we discover we don't have epistemic access to what science wants to describe – the unobservable cells, DNA and the like. We do not have epistemic access, because we can suppose an empirically equivalent world to the realist world, where there are no unobservable entities, or there are entirely different kinds of entity involved. Of course pointing to this sort of underdetermination can lead to problems for the anti-realist. If, like a constructive empiricist, they also wish for observable objects – the kind of objects that might exist while we are not perceiving them – then they need more than the phenomena of perception. In chapter 3 I argued that once we have done some reasoning to establish beliefs regarding external observable entities the same reasoning can get us a range of detectable entities.

One might wonder whether this result is too strong. There is a form of anti-realism that claims that most scientific theories that were once endorsed are now believed to have failed to genuinely refer to what was supposed, by that theory, to exist. The evidence for this claim is garnered from what is known as the pessimistic meta-induction (PMI). I will argue that there are serious problems with at least the strong claims of these arguments (§5.2). I will also argue that PMI type arguments reveal more about the nature of the divide between realism and anti-realism than they provide evidence for one view over the other. Still, some realists have found the PMI forceful and in response, have developed various versions of what we might term 'preservative realism'. The idea behind these accounts is to consider what might be preserved in a theory given that there is something to the PMI. What is preserved can then be explained by what exists. I will consider (in §5.3) an important version of this form of realism offered by John Worrall termed 'structural realism'. Although I am sympathetic to the realist goal, I argue that, as it stands, it is not clear what Worrall's form of realism establishes as real. There are, however, other forms of preservative realism that can account for theory change in science. I consider (in §5.4) what has been termed ‘deployment realism’ offered by Stathis Psillos. While not without its critics, Psillos’ account is an improvement on the structural realism offered by Worrall.

I too will offer a form of preservative realism. Although similar to Psillos’ account, it avoids the usual problems. I will argue that a structure of properties we detect and measure according to the differential approach is preserved or is likely to be preserved as science develops. For the sake of clarity, let us call this form of realism *differential realism*. This term honours Schlesinger’s maxim that the differential method provides ‘the strongest evidence that we are dealing with a ‘real’ property transcending the methods defining it’ [Schlesinger 1958: 305]. In the last section (§5.5) I consider some of the work of Michael Faraday who seems an exemplar of this approach.

5.2 The Pessimistic Meta-Induction

The pessimistic meta-induction is usually associated with Larry Laudan, especially his Laudan [1981], although, interestingly Laudan doesn’t use this term. Even so, there does seem to be some inductive reasoning in some of the arguments he mounts against the realist. For instance, Laudan has pointed to a very long list of respectable scientific theories that failed to refer to their central posits;

> -the crystalline spheres of ancient and medieval astronomy;
> -the humoral theory of medicine;
> -the effluvial theory of static electricity;
> -'catastrophist' geology, with its commitment to a universal (Noachian) deluge;
> -the phlogiston theory of chemistry;
> -the caloric theory of heat;
> -the vibratory theory of heat;
> -the vital force theories of physiology;
> -the electromagnetic aether;
> -the optical aether;
> -the theory of circular inertia;
> -theories of spontaneous generation. [Laudan 1981: 33]

This list, to which we can add examples *ad nauseam,* can include both successful and less successful theories. However, the real interest is in those theories that were successful yet failed to ‘genuinely refer’ to what they supposed existed. For a theory to genuinely refer it need not be successful in what it predicts of such entities. However, there must be ‘entities which ‘approximately fit’ a theory's description of them’ [Laudan 1981: 24]. Still, we might note that judgements of ‘approximate fit’ will need to involve *some* predictive success, otherwise too many theories might genuinely refer. Putting this complication (and others) to one side, the above list seems to provide inductive evidence for the generalisation that all or most scientific theories will fail to 'genuinely refer' to what they take to exist.

This generalisation need not have any implications for scientific realism. Some realists might consider it the case that we will need to eliminate all the erroneous theories in order to discover those theories that are true. There is no reason not to suspect that most theories will fail to refer to what they take to exist. This form of realism can then claim that the fact that most past theories failed to refer supports the proposal that we are increasing the likelihood that we will find the theories that do genuinely refer. Just like looking for a lost object, the more places, where it isn't, that you can eliminate, justifies the belief that you are getting closer to where it is.[1] Even so, the PMI has been used in arguments aimed at other types of realist. For instance, some realists consider that the truth or approximate truth explains the success of science.[2] But if it turns out that, by induction, most successful theories fail to genuinely refer, then most successful theories are false (and false in a way that they can't be approximately true). Thus, approximate truth can't explain, or is poor explanation of success. This is not the only argument available using a PMI. The point of this example is to show up three important features of any such argument. Firstly, there is the evidence upon which a PMI is based, in this instance the above list. Second, there is a generalisation that we infer from the list or whatever we take to be the data. Third, there is an argument using the generalisation. As I will detail below, some attention to these distinctions is quite revealing in understanding what exactly is at stake.

Let us start by considering the data upon which a PMI might be based. Laudan accuses the realist of 'whiggish versions' of the history of science, 'i.e. the ones which recount only those past theories that are referentially similar to the current prevailing ones' [Laudan 1981: 33-4]. Presumably, to avoid this bias we need a fair sample over which we might generalise. Not only do we need theories that are successful but do not (we believe) genuinely refer, we also need to populate our database with those theories which (we believe) do genuinely refer. The list above focuses on some rather grand theories that we now believe failed to refer. However, most theorising in science is done at a much smaller scale. Most modern scientists are likely to study a particular entity or structural feature of the world. It might be a particular bacterium, virus or other causal determinant of a disease, or feature of a DNA strand, or antibiotic, or atomic particle, or geological structure, or distant supernova. One need only consider the sheer number of scientific journals for all the diverse fields of science as evidence of the number current theories and associated unobservable entities available. Nor should this expanded view be limited to pure research. By far the majority of scientists specialise in researching things like electron spin data storage systems, or are involved in developing nano-

[1] Like any induction this example relies on a range of assumptions. For instance, that there is only a finite number of places where the object could be, that the object exists, that you will know when you have found it, and so on.

[2] For an example of this approach see Boyd [1990].

particles for the use in sunscreen, or theorising about the static and dynamic Young's modulus profiles of Lower Cretaceous chalk. Given this broader view, there is in fact an enormous number of theories that are successful (some more than others) and do, according to a realist, genuinely refer.

Given this expanded view, it is not at all clear that the basis for our PMI is well founded. Moreover, there is no reason to think that if the data were to be collected that it would favour a pessimistic generalisation. Laudan proposes, though does not justify, 'that for every highly successful theory in the past of science which we now believe to be a genuinely referring theory, one could find half a dozen once successful theories which we now regard as substantially non-referring' [Laudan 1981: 35]. If it is just a matter of proposing, I propose that for every successful theory that we regard as non-referring there are at least six that are successful and that we believe genuinely refer to what they claim exists. Thus we have a stand off. But the standoff indicates that the basis for the PMI has no more evidential weight than any proposal based on anecdotal evidence.

The next important question is, can the data be collected in a way acceptable to both realists and anti-realists? I doubt that it can. We might note that, for a start, the realist is not interested in the fact that a scientist might *believe* a theory genuinely refers to what exists. Realists that appeal to approximate truth as an explanation for the success of science are not trying to explain how it is that a belief that a theory genuinely refers (or a belief that a theory is approximately true) might be important to success in science. Rather, the issue is the relationship between what exists and what the theory proposes. Thus, these realists will require some criteria for sorting out this relationship in order to gather the data for the meta-induction. But if we have criteria for sorting out what exists and thus its referential relation to a theory, then the game is up for anti-realism. On the other hand, the anti-realist might allow beliefs regarding the referential success of theories, but this fact alone, they will argue, has no implication for realism. But if this is the case, the meta-induction performed will just be an interesting piece of sociology regarding what scientists believe. It is clear, however, if the data regarding beliefs about genuine reference were to favour the thesis that most successful theories are now believed to have failed to genuinely refer, this would somehow count against the realist.

This synopsis might seem a little unfair on the anti-realist. One (common) interpretation of the PMI is that the anti-realist is only proposing that, by the realist's own lights, the truth of the current set of scientific theories, in some sense, entails the falsehood of many (most, all?) of the past theories. So it seems an *internal* PMI is possible. However, it is not entirely clear what is being foisted on the realist here. The anti-realist can't be suggesting the following: Even if we can't know that we have happened upon a true theory, we can know that, if a theory entails the falsehood of most other previous relevant theories, then this creates a data-base of false theories. From this data-base we can generalise, inductively, that most of the theories in this field are false. As the current theory falls within the scope of this generalisation it too is likely to be false. This construal of the 'meta-induction' would take the realist to be a fool. To see how foolish this kind of reasoning is, consider a scientist who thinks it is a legitimate way of formulating a data-base over which we might generalise. He might reason that to increase the inductive evidence for his theory, all he need do is make his theory entail the truth of all preceding theories. This would certainly make the proposed theory complex, *ad hoc*, and maybe even inconsistent, but these worries are beside the point if he has a data-base with which he can calculate the likelihood that his theory is true. There may be problems with a realist interpretation of theory entailing, in some sense, the falsehood of the past relevant theories, but this problem does not create empirical evidence upon which we might formulate a PMI. As it stands it is not obvious why it is a problem that a current theory might entail the falsehood of most past theories, if indeed it does.

More plausibly, what is meant by the internal PMI is that most past successful theories are likely to be false given the *reasons* we consider the current lot likely to be true. These reasons will include evidence from advances made in the development of our instruments of detection and measurement. These advances allow us to falsify past theories, and to the extent that they do, provide evidence that supports our current theories. Although this is a better construal, it is not clear that it supports an internal PMI. Rather than being pessimistic, we might think that, because of these new developments, we have discovered an enormous number of new features or entities of the world. The results of these developments have spawned, we might suppose, an enormous number of new theories: far more in number than those theories that have been falsified as result of such developments. The instruments used to detect and measure oxygen not only falsify phlogiston theory, but they have also allowed us to theorise about all the various oxides that populate our world. The instruments of pathology certainly falsify the humoral theory of medicine, but also continue to uncover and inform us of an enormous range of diseases that we might then research and theorise about. The range of modern telescopes certainly prove that there are no 'crystalline spheres' posited by medieval astronomy, but what wealth of other entities do we now have to theorise about. All these theories are supported, to the degree that they are, by the instruments we use to investigate the world. Apart from this asymmetry, we need not believe, not without justification, that most past theories are falsified (in the loss of genuine referent sense). But even if we think, without justification, that the development of these instruments falsifies most past theories, this does not amount to most theories current and past being false. It is just assumed by the anti-realist (and it seems by many realists) that the number of past successful but falsified theories outnumbers what we take to be the number of current approximately true theories, but why assume that?

There are other interesting responses to the *assumption* that most scientific theories are false. Marc Lange [2002] has pointed out that this assumption by itself need not undermine scientific realism. This is because unprincipled inferences on the basis of past evidence may be the subject of a turnover fallacy. For instance, you might be recently employed as a manager in company. Let us say you discover some employment data for the company that indicates that employee turnover was five fold in the past year. It would be erroneous, however, to infer from this data that if all things remain the same there is a high probability that each of the current lot of employees will quit in the next twelve months. After all, there may be a core group of faithful that will stick with the company through thick and thin. If this is right, there is no in principle reason to worry about the evidence that might, on the face of it, form a PMI. The current theories might be made up of a core group that will see it through to the end.

Juha Saatsi [2005], however, argues that there is a way of formulating the PMI so as to avoid just this defence. Consider the following PMI*:

(1*) Of all the successful theories, current and past, most are taken to be false by the current lights.

(2*) The current theories are essentially no different from the past successful theories with respect to their “observable” properties. (Viz. properties potentially figuring in the realist’s explanatory argument.)

(3*) Success of a current theory is not a reliable indicator of its truth … and there is no other reliable indicator of truth for the current theories.

(4*) Therefore [for all we know] any current successful theory is probably false by statistical reasoning. [Saatsi 2005: 1092]

Key to the success of this argument is that it makes no reference to time dependent properties required for the turnover fallacy to get a hold. It is true that we might try and infer future consequences from 4*, and that these inferences would be susceptible to the fallacy, but this would be to ‘go beyond the validity of … this version of pessimistic induction’ [Saatsi 2005: 1092]. Saatsi thinks that the realist would probably have qualms about premises 2* and 3*, but I think 1* should be considered the main problem. Should a realist really consider that most of the current and past successful theories are false by current lights? Who, we might ask, is shining the light here? It seems the justification for 1* is based on the presumption that the data is available and that it supports the PMI. Given *that* assumption it hardly seems a surprise that 4* follows. It is true that the target of this argument is Lange’s version of the turnover fallacy that assumes something like 1*. My point is that no realist should concede premise 1* without the evidence to support it. If I am right, the anti-realist merely begs the question in proposing something like the PMI*.[3]

There are, however, further lines of attack that do not rely on a generalisation formed on the basis of a large amount of data. Timothy Lyons points out that there is a way of understanding Laudan's evidence as a premise in a *modus tollens* against the realist hypothesis [Lyons 2002]. Understood in this way, the argument requires just one example of a successful yet false theory, and surely the realist would concede this much. He terms this new problem the pessimistic meta-*modus tollens*, and provides us with its general form [Lyons 2002:65]:

> Premise 1: If the realist hypothesis is correct (A) then each successful theory will be true (B)
> Premise 2: We have a list of successful theories that are not true (not B)
> Conclusion: Therefore, the realist hypothesis is false (Not A)

The realist hypothesis in the above argument states that 'theories that enjoy general predictive success are true' [Lyons 2002: 63]. If this is right premise 1 is close to an analytic truth. The only way to deny it would be to show that the realist is not committed to the realist hypothesis as Lyons defines it. Of course, for a realist, the 'realist hypothesis' foisted upon them is too strong. However, Lyons argues that all attempts to weaken it either do not avoid the conclusion, or, alternatively, create more problems for a realist.

The grounds for the realist hypothesis come from the 'no miracle argument'.[4] According to Lyons, the general idea behind the argument is that:

> It would be a miracle were our theories as successful as they are, were they not true; the only possible explanation for the general predictive success of our scientific theories is that they are true [Lyons 2002: 63].

If truth does explain success, *and* for a realist it would be a miracle if there was a theory both false and successful, then they can't avoid the realist hypothesis. Once committed to this hypothesis, there is of course at least one example from a list of many successful but false theories that can be used in the meta-*modus tollens*.

[3] Saatsi also criticises another potential fallacy in the PMI offered by Lewis [2001].
[4] The source of the no miracles argument is credited to Putnam [1978].

Interestingly Lyons also points out that the hypothesis is unacceptable by scientific standards:

> without simply asserting the realist hypothesis to hold, we can establish not a single confirming instance of the hypothesis, not a single instance in which both success and truth obtain.
> [Lyons 2002: 67]

Indeed, according to Lyons, the realists who might try and commit to some kind of *likelihood* of success, given the truth of a theory, also beg the question:

> Without begging the question of realism these modified hypotheses have *no positive, confirming instances*, thus no standard inductive evidence [Lyons 2002: 68; my italics].

Notice that an inductive base for realism, at least according to Lyons, is impossible. But this claim is too strong. Some hypotheses regard what happens in the observable world, surely these can be true because of what they predict? For instance, a cardiac specialist might hypothesise that the symptoms the patient is experiencing indicates a partially blocked artery. This hypothesis could be confirmed by observations made during a surgical procedure. Moreover, some causes may be unobservable at some point in time, but as the reach of our instruments increases, they need not remain unobservable. If this is right, what is hypothesised at one point in time is confirmable at some later point by what we observe. Maybe Lyons can provide a defence against these counter-examples. But if he were to do this, it seems very unlikely that the anti-realists and realists will be able to agree on a criterion such that an empirical basis for a meta-induction might be established.

According to Lyons, despite the realist hypothesis being empirically unfounded, realists are committed to premise 1 because of their commitment to the no miracles argument. Thus a realist, by virtue of what is implicit in the no miracles argument, and the obvious acceptability of premise 2, is either committed to miracles explaining false but successful theories (it would be a miracle if it was the case that a theory is not true and successful), or the no miracles argument is wrong. As we have seen, by the realist's own lights there have been successful but false theories. The more sophisticated realist, according to Lyons, can claim that the more approximately true a theory is, the more successful. But, even if we allow such vagaries as approximate truth there are successful theories that did not genuinely refer to what they took to exist (phlogiston, caloric, various ethers and the like).[5] These theories are not even approximately true, for they

[5] Lyons has an independent argument to show how difficult the concept of approximate truth is for a realist.

did not refer to anything, and the meta-*modus tollens* once again hits hard into the realist flank. Appeals to approximate truth do not avoid the dilemma.

Lyons may have misunderstood the nature of a miracle being appealed to by the realist. In a statistical analysis, scientists look for probable causes for certain correlations. Yet a commitment to a particular cause that explains the correlation need not commit a scientist to miracles in order to account for all the exceptions to the correlation. For a scientific realist then, a 'miracle' is a measurement of the likelihood that the correlation between what is real (according to the theory) and the success of the theory is an accident. Approximate truth explains this correlation, though it need not be the only explanation, or indeed the total explanation. Some of the success of any theory will almost certainly depend on pragmatic considerations, including the initiative and creativity of the scientists involved, the accuracy and sophistication of the technologies available, the financing of the field, and so on. Realism need not posit 'miracles' in order to explain false but successful theories. A sophisticated realist need only argue that:

> The likelihood that the actual explanation for the correlation between scientific theories that genuinely refer and their (predictive) success is a chance correlation is very low. The best explanation involves the approximate truth of the theories involved.[6]

A realist can argue that the miracles argument is merely an allusion to a more sophisticated likelihood of a chance correlation argument common in scientific reasoning. True we have to determine that there is a correlation to explain, and the anti-realist will see any attempt at establishing such a correlation as question begging. But given that begging the question is unavoidable, at least the realist (on the above construal) can avoid the meta-*modus tollens*. Realism avoids the meta-*modus tollens* because premise 2 is compatible with it. Premise 1 can be rejected without affecting the likelihood of a chance correlation argument.[7]

It seems then that the PMI and its variants do not so much highlight a difficulty for the realist explanation for the success of science as much as highlight just how deep the differences are between anti-realism and realism. As Lyons seems to attest, an anti-realist will not accept any theory that is both successful and true. As he sees it, there is no evidence that might count in favour of realism that might not beg the question against the anti-realist. The chance, then, that we might get an unbiased PMI seems

[6] This may not be the only, or indeed the best, construal of a sophisticated no miracles argument. Nor am I endorsing this version.

[7] Musgrave [1988] has considered, in much more detail, how the anti-realist and realist will side up against each other in the face of a well developed 'Ultimate Argument' (miracle argument). His analysis finds them talking past one another given the different goals they give science. I think my analysis provides further insight into these differences.

unlikely. Given this stance, the realist's only option, when looking for criteria with which he might formulate a data-base for a meta-induction, is to beg the question against the anti-realist. Presumably the anti-realist understands that having such criteria means there is no point in performing a meta-induction.

Lyons does provide us with what he considers, and defends, as the most viable position that can account for the success in science, *modest surrealism.* Lyons defines modest surrealism in the following way:

> **MS:** The mechanisms postulated by the theory would, if actual, bring about all relevant phenomena observed, and some yet to be observed, at time t; and these phenomena are brought about by actual mechanisms in the world [Lyons 2002: 78].

According to Lyons MS explains success without committing to the truth, approximate or otherwise, of the theory involved. Given the problems with truth explaining success, this is a distinct virtue. However, we might compare MS with what a strict agnostic is happy to commit to. According to a strict agnostic, there might be no (unobservable) mechanism behind the phenomena. From this perspective we can construe MS as a stronger version of:

> **SA:** The mechanisms postulated by the theory would, if actual, bring about all relevant phenomena observed, and some yet to be observed, at time t; and these phenomena are brought about by actual mechanisms in the world *or not.*

Strict agnosticism considers it a possibility that there are no actual mechanisms that are causally responsible for the phenomena.[8] For a strict agnostic it is a contingent matter whether actual mechanisms are causally responsible for the phenomena; or, alternatively, that there is nothing beyond the observable phenomena. Thus, in order for MS to explain, it has to commit to there being mechanisms in the world that will bring about the phenomena. But the acceptability of SA shows that the existence of actual mechanisms that MA is committed to, in order that it can explain, are truth evaluable. Maybe the advocate of MS would claim that it would be a miracle were there no mechanisms involved in bringing about the observable phenomena. Yet this is precisely the move that MS was designed to guard against. Nor can MS bite the bullet and commit to SA. Unfortunately SA does not explain anything; rather, it is prepared to judge a theory as empirically adequate regardless of the actual mechanisms involved.

[8] I presume here that these actual mechanisms are unobservable. If such mechanisms are, in principle, observable, then a theory that postulates such mechanisms can be confirmed (or falsified).

5.3 Structural Realism

I have argued that a realist need not be concerned about the PMI. There are, however, instances where successful theories ultimately failed to refer to what they supposed existed. This kind of fact still requires an explanation (rather than a response to a PMI). In the next section I argue that attending to what is detected or measured allows us to see what is likely to be preserved as theories develop, while also offering some explanation as to how a successful theory can fail to genuinely refer. (These issues will also be taken up in the next chapter). In this section, I argue that when realists try to take the PMI seriously, (unsurprisingly) the result is a rather thin form of realism. The principal target of this argument is the structural realism developed by John Worrall.

Worrall's development of structural realism claims to be motivated by both the force of the pessimistic meta-induction and our realist intuitions [Worrall 1989]. Firstly, like many, Worrall is drawn to the realistic explanation for the success in science. The real entities of the world make the scientific theories that account for them true, or at least approximately true. Yet in tension with this is the history of successful science, where most (some, many, all?) of the central aspects of our best theories have turned out not to be true. In his paper 'Structural Realism: The Best of Both Possible Worlds', Worrall considers several of the famous shifts in science where theoretical posits are eliminated and a new kind posited. Through a close analysis of the examples of revolutionary shifts in theory, Worrall offers us a way of understanding why a theory was so successful, while the implicit theoretical posits turned out not to exist.

The primary example Worrall chooses to explain his position is the shift from Fresnel's theory of light propagation to Maxwell's. Broadly speaking, this is a shift from an ether theory to a field theory. According to Worrall, 'Light became viewed as a periodic disturbance, *not* in an elastic medium, but in the 'disembodied' electromagnetic field.' Worrall claims that 'one would be hard pressed to cite two things more different than a displacement current, which is what this electromagnetic view makes light, and an elastic vibration through a medium, which is what Fresnel's theory made it' [Worrall 1989: 116]. According to Worrall, even though the different theories made vastly different claims about what exists, there was something in common between the two theories: 'there was continuity or accumulation in the shift [from Fresnel to Maxwell's theory], but the continuity is one of *form* or *structure*, not of content' [Worrall 1989: 117]. This 'continuity' is represented in the form of the equations used in each theory. Crucially, even though the equations themselves mention properties, our structural knowledge is of something different. The form of the equation itself is somehow descriptive of the entities involved. Unfortunately, it is not clear what this description amounts to.

Before we attend to this issue, there seems to be a more serious problem for this account. What might we say of theory change where the form of the equations is not preserved? Newton's theory refers to various properties, force, mass, momentum, acceleration, distance, absolute space, time and the like. According to the structural realist, whatever force, mass, momentum, acceleration, distance, space and time are in Einstein's theory, these properties need not be related, in any qualitative way, to the properties mentioned in Newton's theory. According to Worrall what is preserved in this kind of theory change, as a limiting case, is still a formal aspect of the theories: 'Einstein's equations undeniably go over to Newton's in certain limiting special cases. In this sense, there is "approximate continuity" of structure in this case' [Worrall 1989: 121]. However, what this 'approximate continuity' might amount to is not clear. For anyone who has attempted some analysis of Einstein's theory, formally, Newton's equations are not the same as Einstein's.[9] It is true that there are factors (for example, the Lorenz invariance) that tend to zero at low velocities, and, as a result, these theories can provide similar values for the terms mentioned in each theory. However, this result does not establish that all similar equations have an 'approximate continuity' that is relevant to what exists in the world. There are too many equations that fit this notion of approximate continuity, especially if it doesn't matter what the terms in the equations mean. What is required, it seems, is that each equation has an important relation to the relevant empirical predictions. But even with this empirical grounding, there may be theories that undermine the structure; for instance, Kepler's laws make similar predictions to both Newton's and Einstein's theories. It is true that Kepler's laws don't mention mass or force, but this should not matter for a structural realist. Yet now it is not at all clear – when comparing all three theories – what structure is preserved.

Returning to cases where there is structural similarity in the equations, how might we understand the content of such a description? Obviously, if we are to attribute to something a certain structure, we should be able to describe this structure. We will consider in the next chapter attempts at capturing this structural description using Ramsey sentences. In recent work Worrall [2007] endorses this approach, although I will argue this introduces new problems for his structural realism. In the meantime, we can get a sense of how things might go awry by considering how we might try and describe a structure without providing any further basis to what forms such a structure. Consider for example Newton's second law;

1. $F = MA$

Structural realism says that we cannot commit to what the theory means by the use of the terms *F*, *M* or *A*. So maybe we can describe just the structure by replacing the theoretical terms with theory-neutral mathematical variables *x*, *y* and *z*:

[9] In fact, Michael Redhead [2001: 346-7] argues that it was precisely the *difference* in the structure of Newton's and Einstein's theories that was important in the shift from one theory to the other. One suspects Worrall's form of structural realism can't have a bet both ways here.

2. $x = yz$

Interestingly, there are many theories with just this structure. Consider for instance the following theory:

3. $y_t = \varepsilon_t \, \sigma_t$

Apparently, equation (*3*) has something to do with estimating returns in volatile stock markets.[10] Obviously, the structural realist is not going to claim that the same 'concrete' structure is described in virtue of the structural similarity of both formulas. However, if it isn't the same structure, how does the structural realist rule out this possibility? Certainly not by considering the form of the equations involved, both equations (*1*) and (*3*) have the same form; namely (*2*). A second alternative is to consider the theoretical terms themselves. Unfortunately, what we understand of these terms, according to a structural realist, is irrelevant. A serious shift in our understanding of the unobservable world could change what we mean by these terms. The only option that remains seems to be some appeal to the different empirical domains for each theory. Unfortunately, a fixed empirical domain can support any number of theories, many of which may not be structurally similar. So why, we might ask, prioritise this set of theories over those theories that, although ranging over different empirical domains, are structurally similar? As far as a structural realist can tell, it may be the same structure that explains the success in each empirical domain. It looks as though we need to add something more to an account of structural realism so as to stop it finding structure everywhere.

Yet, even if structural realism is able to restrict its scope to the relevant empirical domain, there remain problems. Without a basis that includes more than structure, structural realism provides very little realism. In order to provide more, a structural realist requires more that structural knowledge. To see this tension, we might consider again the fact that there are an enormous number of theories compatible with a body of evidence. It should be no real surprise, then, that a pair (or many pairs) of theories that range over the evidence could have a similar structure. The ability to find this sort of structural similarity has little or no implication for realism. To garner some realism we might, as Worrall does, draw upon a no miracles argument. It would be a miracle (very unlikely) that the similarity in the form of the *relevant theorie*s is accidental. However, if the no miracle argument does identify a surprising structural similarity between the relevant theories, then, in so far as an account of structural realism might use this fact, it relies on more than just structural similarity and the relevant empirical domain. Assessing what counts as a relevant theory (read successful theory) happens *prior* to any comparison of structural similarity between the relevant theories. But if establishing

[10] σ_t refers to the 'volatility process' of an international stock market, y_t is the 'continuously compounded return on relevant asset' and ε_t is a 'disturbance term'. Taken from Hol [2003: 12].

the relevance of a theory is prior to any assessment of structural similarity, a structural realist cannot be involved in assessing a theory's relevance (success). I conclude that structural realism relies on more than just structural knowledge to explain the success of science.[11]

Part of the problem may have been that the structural realist takes the PMI too seriously. Given that we have undermined the cogency of the PMI, realists need not worry about what a successful theory claims of an unobservable world; at least to the extent that an alleged PMI might motivate such worries. In fact, given the surprising similarity in form between the theories of Maxwell and Fresnel we might now wonder how much difference there is between a 'displacement current' and an 'elastic vibration in a medium'? As it stands these descriptions are rather thin. Maybe, upon further analysis, we might find something preserved beyond the similar form of the equations involved? Taking seriously this proposal is the project of deployment realism. It is something like this form of realism that I wish to defend.

5.4 Improving on Structural Realism

Like Worrall, Stathis Psillos considers actual shifts in theory. However, unlike Worrall, Psillos finds that what is preserved in these transitions is more than just the form of the equations involved. In fact, he argues the scientists themselves are aware of which parts of a theoretical posit are required to get the important empirical results implied by the theory. If this is right, how the scientists use their theories might reveal what will be retained as a scientific field develops. To distinguish what Psillos thinks is preserved, compared to a structural realist, I will term this form of realism 'deployment realism'.[12] Deployment realism considers only those parts of the theory deployed by scientists in order to account for the empirical success of their theory. Although I think Psillos is close to the mark here, deployment realism is susceptible to certain difficult questions regarding the historical evidence used to justify this position. I will argue that a similar form of preservative realism can avoid these problems.[13]

Psillos [1999] gives a very credible account of the development of Fresnel's equations of polarised light propagation, at the interface of two media. Pivotal in the development of Fresnel's equations are those theories regarding the conservation of energy and the well understood wave mechanics of various media. Apparently, the attitude of Fresnel to the wave mechanics used involved a *degree* of agnosticism. Although there was a commitment to something like a medium, there was no explicit commitment to a particular medium. Importantly there was some debate between Fresnel and Poisson on

[11] Psillos [1995: §4.1] considers some further responses to the issues highlighted.
[12] The term comes from Lyons [2002].
[13] Hasok Chang [2003] has termed Psillos' form of realism 'preservative realism'. I think this term is not particularly helpful given that Worrall's structural realism is also interested in what is preserved in a theory.

whether the medium might be fluid-*like* or solid-*like* (the latter seemed more capable of carrying transverse waves necessary for polarised light). The critical point being emphasised is that no commitment was being made to the *complete* nature of the medium, only that the current observations and wave mechanics dictated that there was something *like* a medium. All this is meant to show that Fresnel was not that committed to a *particular* mechanical model.

Psillos [1999] provides a similar analysis of the shift from theories that refer to caloric to thermodynamic theories. Once again, all the characteristics of caloric theory were not given equal weight by the theorists that used them to construct successful empirical laws. One of the central points made by Laudan is that a theory cannot be approximately true if its *central* posits turn out not to exist. Psillos' analysis has shown that in the case of Fresnel's theory and caloric theories, parts of these theoretical entities were *not* central, or if central not in the right sort of way for Laudan's criticism to bite. Just as no particular 'medium' for wave propagation was required for the empirical success of the aether theories, the same goes for caloric theories. In fact, once we locate the parts of the caloric theory that have been preserved, 'caloric theory can be said to be approximately true, despite the referential failure of 'caloric'' [Psillos 1999: 129].

Before we consider some criticisms of this theory, it is important to note how deployment realism might overlap with my account. Elsewhere Psillos [1995] considers the shift from the classical understanding of mass to a Newtonian one. This apparently involved the identification of a 'structural property'. According to Psillos, after Newton, 'mass was understood as having, loosely speaking, a structural property instantiated when a body … [has some force exerted on it] and resists acceleration' [Psillos 1995: 32]. By equating gravitational mass to inertial mass, more structure was established. According to Psillos, knowing what mass is, is knowing what equations it obeys. There may be more to say about an entity or property; for instance, its causal role may not be fully captured in the equations it obeys, but this need not limit our attempts to describe its nature in structural terms. This is obviously very close to what I claim of a structure of properties implicit in functional laws. What I add to this picture is largely methodological. My contribution here is the identification of a method that allows us to test one structure against the other – the differential approach. If I am right about this, we need not wait for another theory to emerge in order to confirm what is likely to be preserved. For instance, because the structure of properties that make up a mass was differentially tested, scientists could commit to this structure. It is true that Einstein's theory changed our understanding of the nature of this structure, but this is just to say that, for a well defined range of contexts we didn't get the description of the structure of properties involved quite right. Einstein showed that there are contexts where our Newtonian understanding of this structure breaks down. This need not threaten the accuracy or relevance of the Newtonian description of the structure in well understood contexts, nor the commitment to the reality of the structure involved or, indeed, the entities that instantiate such structures. This form of preservative realism I term 'differential realism'. We will soon see how it avoids the criticism mounted against deployment realism.

In later work Psillos has laid more emphasis on the practices and beliefs of individual scientists when justifying what might be preserved. Now it is true that it is easy, in hindsight, to find those parts of the theory that were responsible for its empirical success [Lyons 2002: §7]. However, if it can be shown that the scientists themselves were aware of which parts of an unobservable structure were required for a theory's empirical consequences, and which parts to remain agnostic about, then a deployment realist can avoid accusations of historical cherry-picking.

Some critics, however, have questioned whether the evidence does actually support Psillos' position. P. Kyle Stanford [2003] accuses Psillos of being selective in the historical evidence he garners for the agnostic attitude the theorists had towards their various posits [Stanford 2003: 919-20]. Hasok Chang [2003] also finds Psillos somewhat selective in his use of historical evidence specifically to do with caloric theory. In contrast to Psillos, Chang provides evidence to support the view that the material properties of caloric were a central part of the explanation of known phenomena. Apparently, the now discredited aspects of caloric theory were used in the derivation of empirically successful parts of the theory. Even if this disagreement results in a stand off on what the historical evidence supports, deployment realism does not provide an epistemic criterion by which we might arbitrate here.

Differential realism does not suffer this problem. Differential realism is prepared to stick its neck out and provide independent epistemic criteria. We can commit to that structure of properties that we can detect and measure, as outlined in chapters 3 and 4. This may be risky, but it is no riskier than any other empirical claim regarding the observable world, and presumably most anti-realists are prepared to be realists here. Nor is it unrepresentative of actual scientific practice. I will argue in the final section (§5.6) that Michael Faraday was able to use this method to establish a structure of properties for several important theories. Once established, he was able to apply these theories it to investigate further the nature of the empirical world.

Something akin to differential realism is defended by Anjan Chakravartty [1997, 2004]. 'Semi-realism' or 'structuralism', as Chakravartty variously terms his position, makes the distinction between *detection properties* and *auxiliary properties.* Detection properties are 'those [properties] upon which the causal regularities of our detection depend, or in virtue of which these properties are manifested' [Chakravartty 1997: 394]. Detection properties allow us to have knowledge of the hidden world, and if enough are represented, can assure us of the existence of an entity with dispositions to produce the causal properties detected. Auxiliary properties are used to fill out the rest of the story regarding the hidden entities of the world, but because they have not been detected we should remain agnostic. At some later date we can corroborate or falsify claims about auxiliary properties. If falsified, the once hypothesised feature of the entity does not in fact exist. On the other hand, if corroborated, an auxiliary can be converted into a detection property. It should be clear that structuralism is a form of preservative realism.

Chakravartty develops structuralism out of issues that result from the debate between structural realism and entity realism. Each type of realism in their own way has an answer to PMI type arguments mounted by the anti-realist, and each can help each other just where their deficiencies lie.[14] Yet when Chakravartty comes to frame what we ought to commit to – the detection properties – the criterion is developed from the perspective of scientific realism:

> When the realist determines which relations and detection properties are minimally required to interpret the mathematical formalism of a theory, her belief stems from the fact that these are the structures that cannot be denied if one accepts that the theory is reasonably successful in describing parts of the world and their relations to our detectors [Chakravartty 2004: 165].

Although this seems a somewhat vague criterion, Chakravartty does fill it out with some examples. But even if realists see some merit in this approach, surely the anti-realist might hope to comment on what we can commit to here. Parenthetically, Chakravartty mentions that:

> Anti-realists will of course have different views regarding what constitutes an appropriate minimal interpretation and, consequently, what is warranted [Chakravartty 2004: 165].

Unfortunately, he does not go on to consider how a structuralist might respond to this interpretation. Once you admit this view, as Chang [2003: §4] notes, what may be ‘minimally required to interpret the mathematical formalism of a theory’ is just the structure of the data. And this fact, without further development, has no implication for realism.

My position can avoid this result. I have argued that what is measured or detected often goes beyond unaided observation when judged by an anti-realist’s own standards (or at least an important form of anti-realism). The same (differential) method is important when justifying beliefs that are based on what is indicated of the observable world. If the anti-realist were to undermine this method, they do so at the risk of losing perceptual judgements they do want to endorse. Maybe an anti-realist can formulate a

[14] In fact Chakravartty [1997] claims that the two positions entail each other.

theory of perception that can demarcate an epistemic difference in the two types of achievement – detection and observation. However, I am not confident that this can be done, especially since it seems evident that something like the differential method is appealed to in the few examples of perceptual judgements van Fraassen provides.

In the next section I will review an actual case of a theory shifting from being a convenient way of categorising ‘phenomena’ to its use in providing information about what is indicated of the hidden structure of the world. This entitles us to consider as real just what is indicated by the use of our instruments of detection and/or operations of measurement. This form of realism might be more conservative than other forms of realism, yet as a form of epistemic realism it goes beyond what we might claim we observe with our unaided senses. Moreover, it does have something important to say in these revolutionary shifts in theory. Having an independent check is good epistemic practice, even when forming beliefs about the observable world. It is especially important when what we are observing or detecting is something that we are not familiar with. The fact that, in the end, we could not detect or quantify an amount of caloric justifies the belief that these theories failed to refer to just this kind of thing. According to the differential realist, the caloric theorist, and ether theorists for that matter, failed to measure or detect their respective entities because they could not provide two independent ways of detecting or measuring important parts of what they were theorising about.

5.5 What Experiments Indicate: The Case of Michael Faraday

Michael Faraday was one of science’s most prolific experimentalists. His work with electricity, magnetism and electromagnetism resulted in fundamental shifts in the understanding of the kinds of thing that make up the world. Faraday’s work was pivotal in the development of James Clerk Maxwell’s discovery of the electromagnetic theory of light, the development of electro-chemistry and dielectrics, and influenced Ampère’s development of electro-dynamics. Particularly interesting is the fact that Faraday made his discoveries and did much of his research with a minimal use of mathematical theory and a remarkably understated theoretical framework. Furthermore, Faraday was quite strict in distinguishing, in his many papers, proven effects from the deeper more speculative causes. This form of agnosticism is often seen as a fashion or convention of the scientific period. Be that as it may, Faraday took the distinction seriously. I hope to show, however, even within the confines of this conservatism, Faraday provides an interesting example of how the development of new phenomena can indicate, what Faraday might have termed, a ‘real physical truth’.

Faraday did not discover electricity, magnetism or electromagnetism; his brilliance was to show the interrelationship between these kinds. Not only in the design of an experiment that might arbitrate on, or indeed unify various theories, but also his ability to find ways of quantifying, measuring and detecting these new kinds to ever more exacting degrees. Faraday did not so much describe his theory in detail; rather, what he describes in his many notes and articles are the novel experiments and artefacts that he felt displayed what was involved. Some philosophers of science have made much of this ability to use an actual experiment or experimental practice as a 'text' [Baird 1994, Baird & Nordman 1994, Gooding 1991]. Although I am sympathetic to this project, it seems evident that Faraday himself well understood that a display was not enough. Quantification, measurement and detection were the touchstones of a real physical truth. To this end Faraday often used the differential approach.

One of the clearest examples is found in his *Experimental Researches in Electricity* [Faraday 1920]. Here Faraday provides an account of the experiments he conducted on the nature of electricity. In the early nineteenth century, there were at least five kinds of electricity:

- *(i)* Voltaic electricity
- *(ii)* Common electricity
- *(iii)* Magneto-electricity
- *(iv)* Thermo-electricity
- *(v)* Animal electricity

In fact, of these different kinds, the more general nature of electricity was also unknown. So when defining current for his experimental research Faraday's agnosticism is clear:

> By *current*, I mean anything progressive, whether it be a fluid of electricity, or two fluids moving in opposite directions, or merely vibrations, or, speaking more generally, progressive forces [Faraday 1920: §19].

Faraday himself was sympathetic towards the progressive forces view, as this cohered well with his beliefs regarding the nature of magnetism and electromagnetism [Williams 1965: ch. 5]. Although some unification of the kinds of electricity had been attempted, at the time it was not at all clear that unification was even possible. The general approach, it seems, was clear. As the kinds were defined by their source, if it could be found that they all had the same effects, this would be evidence that the same kind was involved. It is a testament to Faraday's remarkable skill as an experimentalist that he was able to elicit the range of known effects from these various sources.

Often success came down to important pragmatic detail. Consider the device Faraday designed for detecting the 'chemical action' of the various kinds of electricity. In figure 1 below, *a* and *b* are tin foil laminates on a glass plate connected to the various sources of electricity via wires *c* and *g*. A solution, usually ‘sulphate of copper’, was placed as a ‘rough line’ *p n* between the electrodes [Faraday 1920: 13]. The bent fine platina wires that make the anode and cathode are pivotal in the design. Firstly, they must be fine enough so that even the smallest amount of precipitate could be detected; but also, because they were easily removed, inspected, cleaned and replaced the test was very efficient. If, for instance, no precipitate was found or the result unclear, further modification could be quickly made and tested. The attention to this sort of design detail allowed Faraday to obtain results where others had failed.

Figure 1.

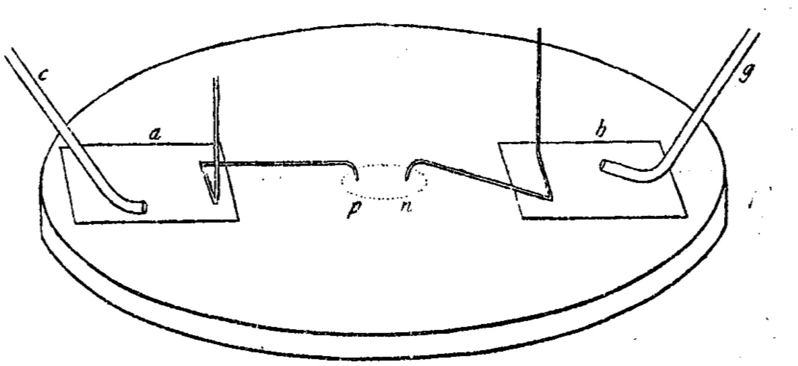

Figure 1. is reproduced from Faraday’s *Experimental Researches in Electricity*. It shows the arrangement used to test for the ‘chemical action’ of the various types of electricity.

After much experimentation Faraday was able to produce the following table [Faraday 1920: 27].

	Physiological effects	Magnetic Deflection	Magnets made	Spark	Heating Power	True Chemical Action	Attraction and Repulsion	Discharge by hot air
1.Voltaic electricity	×	×	×	×	×	×	×	×
2.Common electricity	×	×	×	×	×	×	×	×
3. Magneto-electricity	×	×	×	×	×	×	×	
4. Thermo-electricity	×	×	+	+	+	+		
5. Animal electricity	×	×	×	+	+	×		

The ‘×’, in the above table, represents effects proven by Faraday’s own experiments, while the ‘+’ means results not confirmed by Faraday, but implied by other experiments performed by different scientists. In a footnote to the table, Faraday is confident that the remaining blanks will be filled in, given that: ‘it is a necessary conclusion that they must be possible, since the spark corresponding to them has been procured' [Faraday 1920: 27]. This is a good, if not subtle, indication that the categorisation involves an unacknowledged theory that relates these theoretical concepts. A spark in some sense implies the other phenomena – attraction and repulsion and discharge by hot air.

Regardless of the theoretical motivation, the anti-realist might view Faraday as providing a convenient categorisation of various observed results. Faraday, however, seems to realise that categorisation of the phenomena alone does not settle the matter as to whether or not there is some physical truth behind this categorisation. It was an attempt at quantifying electrical current to which Faraday next turned.

> I next endeavoured to obtain a common measure, or a known relation as to quantity, of the electricity excited by a machine, and that from a voltaic pile; for the purposes of not only confirming their identity …, but also of demonstrating certain general principles …, and creating an extension of the means of investigating and applying the chemical powers of this wonderful and subtle agent. [Faraday 1920: 27]

Voltaic electricity (from a battery or electro-chemical source) and common electricity (static electricity stored in a Leyden jar) were the two most well understood sources of electricity. In order to *measure* a current, Faraday first correlated the number of fully charged Leyden jars with their effect on a suitably calibrated galvanometer. Next, he calibrated the number of voltaic cells needed to produce the equivalent indication on the galvanometer. Having calibrated the operation with two theoretical concepts of electricity he then tested the calibration against an electrochemical effect. To do this, Faraday moistened paper samples with an iodine based solution. The measured amount of iodine that precipitated as a result of the current produced from both sources was the same. This example fits neatly into the model of differentially testing the independently operationalised theories of electricity. From this test, Faraday was able to conclude that:

> for this case of electro-chemical decomposition, and it is probable for all cases, that the *chemical power, like magnetic force…, is in direct proportion to the absolute quantity of electricity* which passes. [Faraday 1920: §113]

Although Faraday states this fact as more evidence for a unified theory of electricity, this important experiment has provided more. Faraday has presented a way of measuring the quantity of electricity from all its sources. With this established method of measuring electrical current, Faraday was able to use these instruments and operations to interrogate his subject matter more deeply. Once this 'physical truth' was established his language was no longer agnostic; rather, it becomes replete with measuring and detecting (or not) an electrical current in innumerable experiments.

So we might well ask: what is it to measure a quantity of electricity, while being agnostic about whether it is a fluid or vibrations or progressive forces? Well certainly strict agnosticism is ruled out. It looks as though Faraday has used a differential approach to establish that some unobservable quantity of electricity is indicated by the use of an instrument of detection or measurement. And what is indicated has a certain structure, albeit thin. As noted above '*chemical power, like magnetic force…, is in direct proportion to the absolute quantity of electricity'.* Of course the strict agnostic could claim that nothing of the sort is indicated; rather, what is provided are more observable phenomena by which we might judge the theory's empirical adequacy. Yet a different kind of method is used to measure an 'absolute quantity of electricity' when compared with those experiments used to unify, even though the same instruments were used with more or less the same observable consequences. By virtue of this method the structure of the data set of a differential test indicates something important. Once differentially established, the experiment is transformed into an instrument of measurement, an instrument that can now be used in further experimentation, but is itself not the subject of the experiment. It is not at all clear that the anti-realist has an explanation for this shift in status of certain experimental practices.

The anti-realist might also be cautioned against being too critical of what we can conclude as a result of this methodology. As we saw in chapter 3, comparing (testing) what is indicated by one's senses can be used to justify a belief about the observable world. It seems Faraday was using a similar method when interpreting what is indicated by the newly established method of measurement. Nor does it make sense to suggest that Faraday was measuring observable phenomena. Rather, Faraday was able to use the various operations and instruments he designed such that they indicate a measurable electric current. There is no suggestion here that what was measured was observable even if some of the consequences of the method were.

Faraday provides a good example of what I am terming differential realism. As noted, differential realism is a form of preservative realism. Crucially, it does not require any comparative analysis between scientific paradigms in order to justify what will be preserved. It makes sense of the fact that scientists themselves understand what is established in virtue of certain type of testing. Differential realism preserves what is observed, detected and measured and those aspects of the theories required for observing, detecting and measuring.

It might seem that what Faraday's example provides is quite thin, and this does not do justice to our intuitions about what is real. Yet we might keep this example in historical perspective, especially given that, at the time, many of the properties involved were considered strange. Still there is no need for a thin structure to stay thin. As we develop more ways of measuring and detecting we can enrich our understanding of a structure of properties. Some of the theories developed by Clerk Maxwell can be seen as enriching the structures discovered by Faraday.

Faraday is probably better known for his work with electromagnetism and magnetism. Here Faraday did have an explicit theory – his famous 'lines of force'. How committed he was to a literal interpretation of these is a vexed question. To complicate matters, his own views evolved, almost to the extent that generous interpretations of his theorising have allowed some to attribute to him the discovery of the field concept.[15] I think it is clear that Faraday did believe, at some stages, in the reality of the lines of force. The important point to note is that in scientific journals and in correspondence, although Faraday pointed to the speculative nature of the lines of force, he also noted real physical truths regarding electromagnetism and magnetism.[16]

As late as 1852, some twenty years after his initial theorising on the subject, he was still qualifying the nature of the lines of force.

> *On the Physical Character of the Lines of Magnetic Force*. NOTE.– the following paper contains so much of a speculative and hypothetical nature, that I have thought it more fitted for the pages of the Philosophical Magazine than those of the Philosophical Transactions …

Regardless, he goes on to explain the importance of the distinction between knowledge and speculation. Referring to his previous use of the term 'lines of force', Faraday suggests that:

> The definition then given had no reference to the physical nature of the force at the place of action, and will apply what ever may be; and this being very thoroughly understood, I am now about to leave the strict line of reasoning for a time and enter upon a few speculations respecting the physical character of the lines of force, and the manner in which they may be supposed to be continued through space [Faraday 1855: §3243].

The definition that Faraday is referring to involved the methods by which electromagnetic and magnetic intensity and their geometric properties could be

[15] Einstein attributes the concept of the field to both Faraday and Maxwell [Einstein 1982].
[16] Despite Faraday's consistent warnings of the speculative nature of his lines of force, Williams interprets him as 'insisting upon the reality of the lines of force' [Williams 1965: 450].

detected.[17] By utilising either the movement of a magnetised needle, the pattern of iron filings made on a surface covering the magnet, or better still, the electricity induced by the movement of a wire crosscutting the ‘lines of force’ (see figures 1, 2 and 3), certain properties of the field could be detected. Importantly, Faraday felt that the induction of electricity was the best method, for it was a procedure that was most amenable to a standard unit [Faraday 1855: §3122, §3177].

Figure 1.

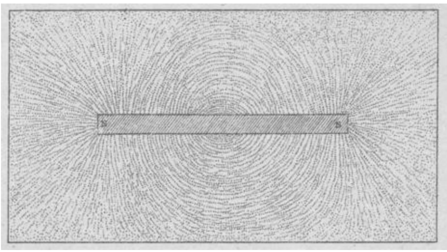

Figure 2.

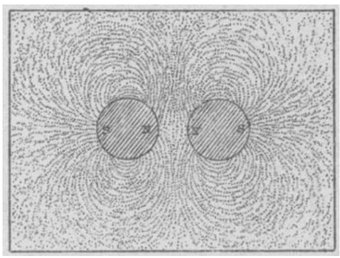

Figures 1 & 2 are typical of many iron filing drawings done by Faraday and used to describe the lines of force of magnetism.[18]

Figure 3.

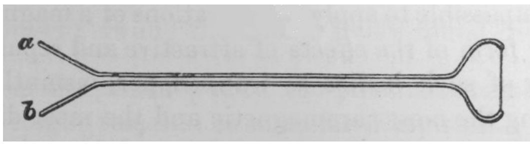

Figure 3 is a tool designed by Faraday to measure the direction and intensity of the lines of force of electromagnetism and magnetism. By connecting it to a galvanometer and moving it such that it ‘cut’ the lines of force, a current proportional to intensity of electromagnetism is generated.[19]

Once again the method used to quantify electromagnetism and magnetism is based on something like the differential approach. And according to Faraday, although we should remain sceptical of these theoretical posits:

> they lead on by deduction and correction, to the discovery of new phaenomena, and so cause an increase and advance of real physical truth, which unlike the hypothesis that led to it, becomes fundamental knowledge not subject to change [Faraday 1855: §3244].

[17] Although obviously related, Faraday kept magnetism and electromagnetism distinct. Ampère tried to unify these kinds as the product of the motions of two electrical fluids moving in opposite direction. Faraday’s testing of Ampère’s ideas led to important developments in electro-dynamic theory [Williams 1985: 99].

[18] These pictures are from Faraday [1855: pl. IV].

[19] This picture is taken from Faraday [1855:§3146].

What this real physical truth amounts to, and its relation to the discovery of ‘new phaenomena’, is of course the issue at stake. So we might briefly step back and consider the nature of Faraday’s speculations. It has been pointed out that Faraday’s ontology was informed by the ideas of R.J. Boscovich, prominent at the time for a concept of matter as a system of forces relative to point locations [Agassi 1971: 80, Williams 1965: ch. 2]. The magnitude of the forces could be described as inversely proportional to the distance from a point. Substance was thus reduced to location and force. In theory then, every point of matter interacts in a network of force with every other point to create a ‘field’ of matter. The interesting problem, Faraday saw, was that the rotational forces implied by the electromagnetic effect seemed to provide something truly unique to the inverse square laws that accounted for electrostatic, magnetic and gravitational forces.

Although Faraday’s theory was meant to offer an explanation for the phenomena, according to Agassi, all Faraday was doing with his lines of force explanation, was trading one incomprehension (action at a distance) for another, rotational line of force [Agassi 1971: 48ff.]. This interpretation is somewhat unfair. Although lines of force might not be fully comprehensible, some explanatory structure is provided in so far as we can appreciate that a force might be transmitted along a line, even if the nature of the line is not entirely comprehensible. As we will see, James Maxwell appreciated this insight, certainly enough to produce one of the most important mathematical interpretations of electromagnetism in the history of the electrical sciences.

Returning to Faraday, probably the best known of his discoveries was the causal relation between magnetism, electromagnetism and electricity (induction). But he also identified that magnetism (and electromagnetism) had an important relationship to light, that it had vectorial properties, and importantly that it had a continuous looped or non-polar nature.[20] These points should be emphasised, for the discovery of this structure to magnetism involved identifying causal relations between unobservable, though detectable, properties. For example, the continuous loop nature of lines of force (the field) through a magnet is not directly observable. It was to Faraday’s credit that he was able to design a series of experiments that indicated just this structure. In virtue of what was indicated of a magnetic field it was discovered that it had a different kind of structure to the non-looped structure of an electrostatic field.[21] Although identifying the similarity and differences between the two kinds of entity, Faraday did not have the theoretical sophistication to provide a functional relationship between the structures that his instruments indicated. It was left up to Maxwell to differentially establish what Faraday’s instrument and experiments indicated.

Maxwell explicitly stated his debt to Faraday, with his theory merely the result of placing the ideas and concepts of Faraday before a mathematical mind. Clearly an

[20] There are poles, but the ‘lines of force’ connecting them are in continuous loops.
[21] For further details see Williams [1965: ch. 10].

understatement, but from a detectionist perspective, this division of epistemic labour provides a neat way of interpreting this important period of scientific development. Maxwell's own words seem to support the idea that he (Maxwell) provided greater theoretical depth to the theories and instruments Faraday developed to measure such structures.

Maxwell interprets Faraday's lines of force as vectors with real magnitudes. After noting the physical force interacting between charged bodies and magnetism, Maxwell points out:

> In this way we might find a line passing through any point in space, such that it represents the direction of the force acting on a positively electrified particle, or elementary north pole, and the reverse direction of the force on the negatively electrified particle or an elementary south pole. Since at every point of space such a direction may be found, if we commence at any point and draw a line so that, as we go along it, its direction at any point shall always coincide with that of the resultant force at that point, this curve will indicate the direction of that force for every point through which it passes, and might be called on that account a *line of force* [Maxwell 1965: 158].

This account, however, is not sufficient for Maxwell. Although accounting for the direction of a force, it does not provide a methodological analogy for calculating the magnitude of the force. To account for this aspect, Maxwell asks us to conceive of the lines of force as tubes of varying section through which a fluid flows. By virtue of this analogy we can conceive that the thinner the tube the faster the flow, and thus proportionally, the greater the force developed. As we observe the tubes radiate out from their source, we can infer that the flow, and analogously the force, is reduced. However, Maxwell constantly warns us to treat this explanatory analogy with a degree of agnosticism. Even so, with the analogy in mind the unification of the forces involved is ready for testing. Maxwell famously goes on to show that light can be understood as involving just this relationship, with Fizeau and Foucault confirming Maxwell's calculated velocity of light [Maxwell 1965: 579ff].

The independently determined speed of light tests the equivalence of theoretical concepts mentioned in Maxwell's equations. Something of the structure implicit in Maxwell's equations was independently understood in the less sophisticated theoretical work of Faraday. This was an enriching of our understanding of the structures involved. We should not forget that this enriching goes hand in hand with the instruments and experiments used to indicate just these structures. Moreover, because something like the differential method was involved we are justified in believing our claims regarding a structure of properties will be preserved. Importantly, we don't have to wait for the next theoretical revolution in order to justify the belief that this structure will be preserved. Although there is some epistemic risk when interpreting what is indicated by our instruments of measurement and detection, the differential approach significantly reduces this risk.

Faraday and Maxwell provide good examples of the differential approach used in scientific practice. Scientists, it seems, understand well enough the need for an independent check on what they claim exists in the world. Once we appreciate how such a check might be achieved, the differential realist has a way of assessing what will be preserved in a theory. If this is right, differential realism provides an antidote to what might remain of the PMI type arguments. What it is for parts of a scientific theory to genuinely refer will be explored further in the next chapter. However, it seems reasonable to suppose that, given the differential approach, if we can detect or measure important structures we have a justification for the belief that the relevant parts of the theory genuinely refer to what is measured and detected. If what we detect or measure is not consistent (according to the theory) then we can falsify the theory and claim that it does not genuinely refer. An additional advantage of this approach to scientific realism is that it provides an explanation for the shift in our epistemic attitude to a theory. As we saw above, Faraday provided a good example of how the treatment of a theory can shift from something to be tested and maybe falsified, to something that can be used to investigate further the hidden features of the world. Differential realism provides a good explanation, and more importantly justification, for this shift in attitude.

Finally, I should mention that, although differential realism is concerned with how we might warrant what is likely to be preserved of a theory, I should not be read as arguing against the more ambitious forms of scientific realism. For instance, differential realism need not undermine other forms of realism that argue for some supportive relationship between the evidence and the more hypothetical parts of the relevant theory. However, differential realism does identify an important distinction between the secure and hypothetical parts of a theory that any account will need to accommodate.

Chapter 6: Ramsey Sentence Realism

6.1 Introduction

In the previous chapter, I assessed what Stathis Psillos [2001] has termed the '*downward path'* to structural realism. The downward path begins by identifying a structure present in important parts of several theories. The same structure is then said to be explained by some real feature in the world. In virtue of the downward path, structural realists are able to distil a structure from the theories that somehow reveals something about the *real* structure of the world. Unfortunately, if it is only a formal structure that is identified, then it is not clear what real unobservable thing must exist to account for the similar form. As we saw in the previous chapter, many theories across many different fields share the same form, yet we might hope to resist the conclusion that the same concrete structure explains this similar form. It is not clear that the structural realists can solve this issue without some non-structural theoretical commitment. At the other end of the spectrum we have also noticed that (pure) entity realism might allow us to refer to something because of our interventions. Unfortunately, unless we commit to some theory realism, just by virtue of such interventions we can say very little about the unobservable something. Some forms of realism (deployment realism, semirealism, differential realism) defend the proposal that there is some descriptive content of a theory preserved or likely to be preserved, and that we are justified in claiming that this description is of something real.

As a minimum, I have identified the need for a structure of properties to which our laws and theories might refer and describe because of what we have detected and measured. Entities will instantiate these structures, though the entity itself need not be considered the only 'unit of reality' described in a theory. Structures of properties can be considered real even if the complete nature of what is involved is unknown. As we establish more structure, in virtue of what we measure and detect, we enrich our understanding of the entities that are a part of, or instantiate, these structures of properties. On the other hand, if the properties of a structure of properties can't be measured or detected using the differential method, we have good reason to believe that the theory, to this extent, fails to refer to what it supposes is the relevant structure of properties. I have termed this form of preservative realism, differential realism.

There is, however, what Psillos terms an 'upward path' to structural realism. According to the upward path we are able to glean off from what is observable something of the structure of the hidden entities involved. Thus, it seems, we are able to provide some content for a theory that accounts for the observables. Although Worrall originally took the downward path, he now endorses its upward route. More accurately, he sees the two approaches as more or less equivalent. An important tool in the analysis of the upward path is the Ramsey sentence – named after an approach developed by Frank Ramsey

[1929/1990]. However, something similar was proposed by Bertrand Russell [1927], and developed further in Grover Maxwell [1970a, 1970b]. The Ramsey sentence programme has also been refined by Carnap [1966] and then Lewis [1983] and is currently having something of a renaissance in sorting out issues in the philosophy of science.

In this chapter I intend to review some of the uses of the Ramsey sentence approach. Overall, the conclusion is that the use of the Ramsey sentence in these attempts is not very helpful, or needs to be supplemented to get the right result. The first part of this chapter (§6.2) reviews how the Ramsey sentence approach might help a structural realist. I find that to avoid problems created by unguided Ramsification, other criteria need to be introduced. But then it is not clear what explanatory work Ramsification is doing over and above what is introduced to avoid these problems. The second part of the chapter (§6.3) considers the use of the Ramsey sentence in sorting out issues of approximate truth and referential success in the realism debate. I argue that something like differential realism is required such that the Ramsey sentence approach can explain a theory's empirical success. If this is right, then this requirement imports issues that the Ramsey sentence approach is supposed to avoid.

6.2 The Upward Path to Structure

Bertrand Russell thought that the structure of the hidden world was revealed in the structure of our perception of it. His approach, if adopted by a structural realist, is thought to face serious problems.[1] Just how serious these problems are I leave to others. (Worrall does not believe them to be as serious as the critics venture [Worrall 2007, Zahar & Worrall 2001].[2]) Even so, a brief review of these problems will help distinguish them from the issues I wish to highlight. In this section I will argue that the Ramsey sentence approach contributes little to the understanding of what it is for two or more theories to share a structure. This is because we either need some other epistemic criterion in order to avoid a problem that an unqualified use of the Ramsey sentence approach generates; or, alternatively, the approach only identifies what is already known to be true. In fact these alternatives form part of a deeper problem for the strict structural realist who regards the Ramsey sentence as having something of a dual purpose. When we consider more viable Ramsey sentence approaches, we discover that we have to have some way of 'fixing' or 'grounding' what we predicate of the entities involved. I suggest that this fixing is to do with measuring and detecting. But whatever the right account is, if it is the method of fixing the relationship between several

[1] For a critique of Russell, see Newman [1928]. Or a modern version, Demopoulos & Friedman [1985]. Both Ladyman [1997] and Psillos [2001] draw out similar implications for Maxwell's project. For a more general and formal treatment of the problem, see Ketland [2005]. For an account of how Ramsey himself might have construed these matters, see Psillos [2009: ch. 9].

[2] See also Votsis [2003] for a defence of structural realism. Joseph Melia and Juha Saatsi [2006] also argue that the project can be repaired so as to avoid these issues.

predicates that establishes what we can commit to as real, then, once again, it is not clear what the Ramsification of a theory does for structural realism.

According to Russell [1927: ch. 20], notwithstanding how any particular theory might account for the organisation of our percepts, some structure, attributable to these percepts, reflected, or was isomorphic with, an unobservable structure. Russell develops this form of structuralism as the causal theory of perception. He contrasts the causal theory with the phenomenological theory of perception and solipsism, and concludes that the causal theory is the better explanation. Russell next considers what we can 'objectively' infer from a 'group of percepts' [Russell 1927: 222ff.]. It is important to note that, for Russell, a group of percepts are the causal effects of the observable world. Given this view, two or more persons can, from different perspectives, regard their percepts belonging to a single group caused (by hypothesis) by a single object. What we count as *observable* is in fact an explanation for these effects. For Russell then, observables are theoretical posits. Although on the face of it this seems to run against the 'common sense' view, Russell claims his account is the scientific understanding of perception. Accordingly, while we can gain knowledge of the quality and structure of the percepts, we only have structural knowledge of their causes. At best, the structure of our percepts preserves, or is isomorphic with, the entities that caused them.

Unfortunately, in so far as we might call this knowledge of the unobservable world, it is not clear how a *formal* analysis of a group of percepts might reveal this. Russell was aware of this problem, but thought that something of the structure of the unobservable world could still be derived from such an analysis. After all, surely the world itself had something to do with the way such percepts were organised. Yet, if a group of percepts is just a set of percepts then, as M.H. Newman [1928] first pointed out, the only empirical structure *formally* preserved here is cardinality, as the only formal structure of a *set* of percepts *is* cardinality.[3] This weakness makes the method too strong; for if the structure of a set of percepts has no structure apart from cardinality, then *any* set of the right cardinality is isomorphic with the formal structure of a set of percepts. As a result, the truth of a theory that accounts for the percepts is *ensured* by formal logic ranging over any set of the right cardinality (for more detail, see the references in note 1).

One idea that immediately suggests itself is to posit some basic relations between percepts and the hidden entities that account for the percepts: fundamental correlations, interpretations or even causal relations.[4] But this, apparently, is to move away from *strict* structural realism, since we are now claiming that we know something apart from

[3] It is true that Russell does consider the role of a 'set of percepts' [Russell 1927: 231] in relation to the laws of nature, but in so far as he considers the structure of perception he makes use of the concept of a 'group of percepts' Russell [1927: 222ff.]. Given that Russell did see the criticism made by Newman [1928] as valid, presumably he did think the two kinds were importantly related. However, it remains an open question as to how the two kinds are related.

[4] Russell suggests this option in response to Newman, but leaves it undeveloped [Demopoulos & Friedman 1985: esp. 632].

the formal structure of an unobservable. It seems, then, that this approach is caught on a fork: either it is so trivial to be true but says nothing, or, in so far as it has to posit non-observational, non-formal relations to a structure, then *strict* structural realism is false.[5]

The initial plausibility of Russell's causal theory of perception is undermined if what we consider as objective in perception amounts to a *set* of percepts (see note 3 above). However, when we observe something, what we observe is not a set of percepts, or not just a set of percepts. For a start a set *qua* set is a mathematical or logical device. So construing a perception as a set of percepts looks like it is destroying the very structure the causal theory of perception explains. It is not surprising that trivial results ensue. The trouble is that Russell, after justifying a causal theory of perception on the basis of explanatory power, then tries to construe these causal relations formally, thus allowing Newman's critique to bite.[6]

This problem has not stopped others from developing similar accounts. Grover Maxwell, for instance, further develops the strict structural realist approach [Maxwell 1970a, Maxwell 1970b]. His account is instructive, as it exposes a tension in the Ramsey sentence methodology that I wish to bring to light. So let us first explicate what it is to Ramsify a theory according to Maxwell. Consider the following:

1. $\forall x((Ax \,\&\, Bx) \supset \exists y Cy)$

Maxwell asks us to interpret (*1*) above in the following way: let *Ax* stand for 'x is a radium atom', *Bx* stand for 'x radioactively decays', and *Cy* stand for 'y is a click from an appropriately located Geiger counter'. Both *A* and *B* refer to theoretical properties while *C* is claimed as observational. A structural realist would not want to commit to *A* and *B,* since a significant ontological shift in science might remove what we refer to as a radium atom and/or the process of radioactive decay. There is, however, some structure in the relationship between the unobservable and the observation – the click. And it is here that the structural realist doffs his cap to Frank Ramsey. The basic idea behind Ramsey's proposal, despite its formal presentation, is quite straightforward. A theoretical term in a theory that refers to a property or kind of entity can be replaced with a 'free variable of the appropriate type' [Maxwell 1970b: 187]. Rather than claim that 'there is a radium atom such that …' we can claim instead that 'there exists a ψ such that …'. Here ψ is not just another name or synonym for radium; that would do no work at all. Rather, construing ψ as a variable allows us to loosen the link between the meaning of the theoretical term and what it refers to. Indeed, by quantifying, the

[5] Votsis [2003] provides a defence against this criticism by utilising some Quinean equipment. By appealing to background theory, we do in fact have a way of selecting from the many possible structures consistent with the data. Moreover, given this ontological relativity we need not commit to any basic properties of this ontology. Well this is fine as far as it goes, but we need some account of how or why this implies any form of realism.

[6] Van Fraassen has recently considered Russell's structuralism and how a constructive empiricist might avoid the problems that beset that programme [van Fraassen 2008: Pt. 3].

Ramsey sentence picks out *whatever* it is that accounts for the observation. Our knowledge of a theoretical posit, and thus what is meant by a theoretical term's use, is provided indirectly. What ψ means is captured in the logical form of, and thus relation to, the relevant observation terms (and other variables) within the Ramsey sentence. Thus our knowledge of the unobservables, according to Maxwell, is by description rather than acquaintance.[7] So, returning to the above example, Maxwell Ramsifies the above statement (*1*) to reveal the minimal structure we can commit to:

2. $\exists\psi\exists\phi\forall x((\psi x \ \& \ \phi x) \supset \exists yCy)$

We can read (*2*) above as saying 'there are some properties ψ and ϕ such that if anything has those properties, then an appropriately placed Geiger counter will click'. To be sure (2) is a very short Ramsey sentence, and consequently it does not provide much knowledge of the theoretical posits involved. Maybe if we quantified every use of the term 'radium atom' and 'radioactive decay' in our current science we could know more of what we mean by whatever fills the role of ψ and ϕ? But before we grant this, let us consider what is being claimed in the above Ramsey sentence. We might consider the term 'appropriately placed' in the observational predicate above. What makes the Geiger counter appropriately placed, it seems, has something to do with the nature of the antecedent of the conditional. That is, the Geiger counter is placed *near* a radioactively decaying radium atom. But how do we predicate of some unobservable this relational property *near*? The term 'appropriately placed' implies that we know that whatever the unobservable is, it is near the Geiger counter. But some future theory might propose that the unobservable posit is the kind of thing that is non-locatable. What counts as 'appropriately placed' now seems to be an important theoretical consideration.

Once we have identified all the theoretical terms, the Ramsey sentence does not seem to be saying much: something like, for any *x* it either does not have two particular properties or a Geiger counter will click.[8] Thus stated, whatever the Ramsey sentence captures of the meaning of a theoretical term, it seems rather thin. It does not say much because it looks trivially true. As a material conditional, if it is true that there is a clicking Geiger counter the conditional is true.[9] We might try to reduce this problem by strengthening the scope of the universal to something like:[10]

[7] This knowledge is of the 'structural characteristics' of the theoretical posits; where, 'structural properties are always of a higher logical type: they are properties of properties, properties of properties of properties, etc.' Maxwell [1970b: 188].

[8] Formally, Maxwell's Ramsey sentence is equivalent to $\exists\psi\exists\phi\ \forall x(\sim(\psi x \ \& \ \phi x) \vee \exists yCy)$.

[9] Achinstein [1968: 85] has pointed out that there is another way of reading conditionals like (*2*) to make them vacuous. We might notice that (*2*) is satisfied if we can find two properties that nothing has. Candidates for the two properties nothing has might be, being a perpetual motion machine and being a gas at ten billion degrees. This makes our Ramsey sentence trivially true, or true for the wrong reasons. Hintikka [1998] claims that Achinstein misses the mark on formulations like (*2*), as it can be seen that formulations like (*1*) also have no observable consequences; after all (*1*) is just a conditional. But in so far as (*1*) is empirically adequate it is asserting the truth of the observation statements it is supposed to

3. $\forall x((Ax \;\&\; Bx) \supset C'x)$

Ramsifying *3* we get:

4. $\exists\psi\exists\phi\forall x((\psi x \;\&\; \phi x) \supset C'x)$

In order that *3* and *4* make some sense, we will have to modify *Cx* to say something like, '*x makes* the 'appropriately placed' Geiger counter click' (call this $C'x$). But this is no good, for it either attributes to the unobservable entity an observable property C', which on the face of it seems paradoxical; or, alternatively, C' is a causal relation predicated between some *x* and something observable. The second approach seems more promising, yet causal relations, we might suppose, are the very thing we should theorise about. If this is right, causal relations should be Ramsified leaving the observable predicates as something upon which we might ground the theory. In fact, John Winnie [1967: 228n] credits Maxwell for pointing out some of these complications to him in the development of his influential paper, 'The Implicit Definition of Theoretical Terms' [Winnie 1967]. Winnie suggests there may be a way of predicating causal properties of some *x* so as to avoid these issues, but at the time the programme was undeveloped. David Lewis took up Winnie's challenge; however, as we will soon see, if we take up Lewis' solution, we get much more than strict structural realism.

We will consider Lewis' approach shortly. However, even prior to attending to this problem there is a deeper problem regarding the use of the Ramsey sentence method by the strict structural realist. Given that there are different explanations for the same observations, surely the meanings of the terms used in each explanation are different? If the terms mean different things, then Ramsification, in so far as it is involved in understanding what the terms mean, should respect this difference. Yet according to the structural realist it is precisely because the Ramsey sentence only preserves structure and observational content that they consider it a viable tool for what can be said as *similar* when comparing *different* theories.

We can see this problem explicitly by developing Maxwell's example. There seems to be no loss (at least for the sake of this illustration) in combining all the theoretical

---

explain. So neither the theory nor the Ramsey sentence is observationally trivial – they both assume the observation statements. Yet even with this correction, surely the truth of the theoretical explanation of (*2*) is true for the wrong reasons?

[10] Joseph Melia and Juha Saatsi [2006] think that including intensional relations saves structural realism and thus avoids the trivialisation problems. This approach would involve developing second order relations between properties rather than relations between individuals. These relations could be nomological only or natural necessities. This allows the structural realist to identify relations between properties rather than individuals. For example the property *being water* can be 'necessarily correlated' with the property *being H_2O*. This is all grist for my mill. The contribution of differential realism is to add some epistemic criteria for such judgements.

properties into one; where *Rx* now says '*x* is an appropriately placed radioactively decaying radium atom'. This leaves *Cy* to say '*y* is a click from a Geiger counter'.

5. $\forall x(Rx \supset \exists yCy)$

Let there be a competing theory that explains the click with an appropriately placed micro-butterfly. Although this presentation is sparse the explanation that relates the theory to the empirical results can be 'enriched' so that it is quite specific about what it means to be a micro-butterfly. So, for instance, the butterfly theory might claim that the Geiger counter is actually counting a clicking sound of the beating wings of the micro-butterfly. As with the radioactive decay story we can combine the theoretical properties mentioned in the micro-butterfly explanation into one predicate *M*.

6. $\forall x(Mx \supset \exists yCy)$

Let the theories now define themselves as mutually incompatible. For instance, it cannot be the case that half the clicks counted by a Geiger counter are from appropriately placed micro-butterflies and half from radioactively decaying radium atoms. It is not the case that there are both appropriately placed micro-butterflies and appropriately placed radioactively decaying radium atoms.

7. $\sim(\exists xMx \;\&\; \exists xRx)$

Thus we might reformulate the theories to get:

5′. $\forall x(\sim Mx \;\&\; (Rx \supset \exists yCy))$

6′. $\forall x(\sim Rx \;\&\; (Mx \supset \exists yCy))$

When we Ramsify, according to a strict structural realism, we preserve what can be truly said of empirically adequate theories that range over the same observational content. Hence we get:

5″. $\exists\psi\exists\phi\forall x(\sim\psi x \;\&\; (\phi x \supset \exists yCy))$
6″. $\exists\phi\exists\psi\forall x(\sim\phi x \;\&\; (\psi x \supset \exists yCy))$

Yet it might seem that, insofar as we have eliminated the incompatibility between 5′ and 6′ and claim that *5″* and *6″* are equivalent, something of the meaning of the theoretical terms involved is lost. It is hard to see how, but if we do maintain that the incompatibility of the terms expressed in 5′ and 6′ lives on in 5″ and 6″ then it is not clear what additional insight the Ramsey sentences of 5′ and 6′ provide for structural realism.[11] We can compare theories as much as their Ramsey sentences if they do not eliminate incompatibilities.

A structural realist might claim that this way of construing matters just begs the question. Yet sometimes theories have built into their structure the falsity of competing theories, despite empirical equivalence. For instance, some interpretations of quantum phenomena apparently entail the impossibility of 'hidden variables' [Baggott 2004: 144]. Maybe a structural realist could just deny that the two theories 5′ and 6′ say different things in virtue of their equivalent Ramsey sentences 5″ and 6″. Different theories that say different things regarding the unobservable world will have different structures captured in their Ramsey sentence, even if they have the same observational content. Unfortunately, on this assumption, *different* theories that are empirically equivalent will have a different structure, and thus will not, indeed should not, be structurally similar, even after Ramsification. On the other hand, if the two Ramsey sentences of two theories are found to be equivalent, then the two theories are not different theories (in anything other than the terminology they use), and thus it is hardly surprising that we might consider them structurally similar. What, we might ask, does Ramsification add to our understanding of the structure preserved in a range of theories?[12]

It might seem I am making a straw-man out of structural realism, but Worrall claims that nothing of the cognitive content of a theoretical term is lost by the process of Ramsification.

> Quantifying over erstwhile theoretical predicates removes them linguistically, but the theoretical terms live on within the Ramsey sentence via the structure that they impose on the observational content …

And further on:

> [W]e fool ourselves if we think that we have any independent grip on what the [extensions of the theoretical terms] are aside from what ever it is that satisfy the Ramsey sentence' [Worrall 2007: 152][13]

Even if the above problems were overcome, when we expand what might be Ramsified to consider all theories that are compatible with the relevant observational content, the

---

[11] It is unlikely that Worrall would endorse trying to preserve this sort of incompatibility [Worrall & Zahar 2001: 248].

[12] Psillos [2009] has found further problems with Worrall's use of the Ramsey sentence to support structural realism.

[13] See also Worrall's discussion on the cognitive content captured in a Ramsey sentence and its relation to a causal theory of reference [Worrall 2007: 148].

comparable structure must become too thin to carry any realist implications. If a Ramsey sentence approach reveals structures amongst *all* theories that range over the same observational content then it seems the structural constraint runs in the reverse direction. It is the observational content that imposes a structure on the theories. But this is precisely the point that a constructive empiricist would hope to make, claiming it is merely a pragmatic consideration by which we choose a theory that fits the observational content.

We have already considered something like this problem in the previous chapter. We might ask, does the Ramsey sentence approach offer any help? As mentioned, Worrall claims one of the main motives for being a realist is based on something like the 'no miracles' argument. In fact he thinks that the concept of simplicity or the notion of explanatory power that such arguments appeal to is required to set a 'default position' when considering empirically adequate theories [Worrall 2007: 147]. This manoeuvre allows us to rule out *ad hoc* theories that range over the same evidence. However, if it is the 'no miracles' argument, or some feature of it, that is doing epistemic work in his structural realism then it is now not clear what additional work Ramsification is doing. To see why, let us return to the original presentation of structural realism.

The power of Worrall's [1989] original approach turned on our ability to note a structural similarity between functional laws of different theories; for instance, those developed by Fresnel and Maxwell. Originally it was a surprising similarity that was noted in the form of the equations that motivated the idea that this structure might represent something in reality. Given that this similarity was not an accident, or a 'miracle', the best explanation is that there is something real that accounts for the similarity. One of the equations that Worrall uses to point out the structural feature was '$R/I=tan(i-r)/tan(i+r)$'. Now it is not at all clear how we might construct a Ramsey sentence of this equation. For a start we might ask, do we need to 'conditionalise' the functional law in relation to actual observational content? In its functional form, the equation does not refer to any actual observations. Assuming these issues can be resolved, we might wonder how we got to the point where we can compare the structural similarity of just these particular equations. As mentioned, the Ramsey sentence approach can range over any theory consistent with the same observational content. In order to get the right theories to Ramsify we need to eliminate all the *ad hoc* theories. Presumably it is not a miracle that two very similar equations are to be compared for structural similarity. However, we might well ask, in so far as a strict structural realist is able to get this far with the help of the 'no miracles' argument, what then does Ramsification of the theories add?

For a strict structural realist, it seems, one purpose of Ramsification is to capture in the structure of the Ramsey sentence the meaning of a theory's theoretical terms. A second purpose is to capture what can be truly said of two or more *competing* theories ranging over the same empirical domain. In so far as the strict structural realist tries to achieve the second goal he must eliminate important differences between the individual theories. Thus, prior to Ramsification we must either identify what structure is to be preserved, or, alternatively, we have to Ramsify in such a way that competing theories that range over the same empirical domain can say the same thing. The first option is achieved with the help of the 'no miracles' argument that excludes those competing theories that might corrode the identified structure. But then we might ask, what additional insight is Ramsification providing? The second option is in stark contrast to one of the goals of Ramsification: to capture, in the structure of the Ramsey sentence, the meaning of the theoretical terms. If the meaning of a theoretical term does live on in its Ramsey sentence, without loss, then so does its competitive streak. If this is right we will end up with several Ramsey sentences that will need to be compared for structural similarity. But once again, it is not clear what additional insight Ramsification provides over and above 'pointing', as Worrall [1989] originally did, to the similarity between the original theories.[14] If the Ramsey sentences of different theories are equivalent, or can be made equivalent, it is not clear why we should consider the theories different in anything other than the terminology chosen, and thus it is hardly surprising that they are structurally similar.

Maybe this just shows up a problem with the Ramsey sentence approach chosen. Maybe we can allow some flexibility in developing a Ramsey sentence methodology to structural realism, while staying within the spirit of the programme. We cut short our analysis of the problematic nature of predication to consider the problems associated with the use of the Ramsey sentence by the strict structural realist. Maybe both sets of problems can be solved at once. According to some, the Ramsey sentence approach might provide a viable account of realism if we can introduce well 'understood' *mixed* predicates capable of attributing the same property to either observable or unobservable entities [Cruse & Papineau 2002, Cruse 2005: 561, Melia & Saatsi 2006: 567,]. Good candidates for mixed predicates are '__ is larger than __', or '__ is inside __', but can also be non-relational predicates such as, '__ has a mass', or '__ has a location', or even causal relations '__ because __', or '__ tightens __'.

Pivotal in the development of this programme is the work of David Lewis [1972, 1983]. The title of his principal work on the Ramsey sentence is, 'How to Define Theoretical Terms'. Straight up, we might remind ourselves that science does not just *define* its theoretical terms; it is also in the business of making empirical discoveries. Theories and what they predicate of the hidden world are meant to refer to the properties and entities in the world. Scientists then, when developing a theory, are doing several things: they are defining, while also describing and explaining discoveries. Lewis acknowledges Carnap's important development of the Ramsey sentence in providing a solution to this problem. Carnap's idea here was to conditionalise the Ramsey sentence

[14] As I argued in chapter 5, it is also unclear what pointing to *just* this structure establishes for a realist.

so as to make, what he calls, an A-postulate. The general form of an A-postulate ('A' for analytic) is as follows:

8. $\quad {}^{R}TC \supset TC$

The ${}^{R}TC$ of (*8*) just refers to a Ramsey sentence like (*2*) above. The consequent *TC* is the un-Ramsified version of this sentence; which is a conjunction of all sentences in which the theoretical terms we are trying to define occur.[15] Roughly, *TC* is our theory and ${}^{R}TC$ is the Ramsification of the theory. This move allows us to define our theoretical terms within *TC* independent of the empirical content captured in ${}^{R}TC$. We can then go about trying to discover if ${}^{R}TC$ is in fact the case. If it is realised, the theoretical terms in *TC* have an interpretation (what *TC* says is true of some thing). If ${}^{R}TC$ is not the case, then (*8*), the A-postulate, need not be meaningless even though there is no realisation for *TC*.[16]

Working with the insights of Carnap and Ramsey, Lewis then divides our language into *T*-terms and *O*-terms. However, unlike his predecessors, his distinction is not between observational terms and theoretical terms. Lewis claims he cannot make sense of such a distinction. All predicates are *mixed*, in the sense suggested above. The distinction is only that *O*-terms should be considered 'well understood'. Apparently '*other'*, '*original'* or '*old'* terms of a theory correspond well those terms that are well understood.

In contrast to *O*-terms there are *T*-terms ($\tau_1 \ldots \tau_n$) that are new, or introduced, or less well understood. The various τ_i's *name* 'entities of any kind: individuals, species, states, properties, substances, magnitudes, classes, relations, or what not' [Lewis 1983: 80]. We can, it seems, make a name out of any predicate. For example, in the case of the predicate '___ is a substance' we can redescribe the predicate as a property, terming it 'substance-hood' such that we can claim '___ has substance-hood'. We can represent 'substance-hood' as a *T*-term with the constant τ_1 say. The predicate '___ has ___' is apparently well understood, and thus is free to be an *O*-predicate. If 'substance-hood' becomes well understood, it is free to be an *O*-term. Lewis claims that there is no reason not to give *O*-terms a *fixed* interpretation [Lewis 1983: 84]. Given the above, a postulate of our theory *T* can be formally presented thus;

9. $\quad T[\tau_1 \ldots \tau_n]$

It should be noted that although (9) includes *T*-terms, it also includes *O*-terms and predicates. Formulating the Ramsey sentence then is simple enough:

[15] Carnap [1966:232ff.] has more to say on defining the terms with C-postulates ('C' for consequence).
[16] Carnap develops this approach in a range of publications. Psillos [2000] has done much to uncover the depth and development of Carnap's ideas here.

10. $\exists x_1 \ldots \exists x_n\ T[x_1 \ldots x_n]$

We can conditionalise (9) and (10) to produce what Lewis calls the Carnap sentence ((8) above):

11. $\exists x_1 \ldots \exists x_n\ T[x_1 \ldots x_n] \supset T[\tau_1 \ldots \tau_n]$

Lewis goes on to provide the conditions for unique reference for theoretical terms. What seems interesting, however, is the use of fixed mixed predicates that might allow some non-observational content to be retained through a theory change. Worrall explicitly rejects Lewis' approach on the grounds that he finds the uniqueness criteria 'unconvincing' [Worrall 2007: 149n.]. Even if unconvincing, the use of fixed mixed terms might provide some analysis of terms describing causal relations or powers or being appropriately placed near a Geiger counter.

It is important to emphasise what might be involved in grounding mixed predicates of some unnamed *x* such that they might be retained through a theory change. Maybe we understand the *O*-term in the *O*-predicate '__ has a rest mass of 9.11×10^{-31} kilograms' independent of some specific unobservable thing and whatever comes and goes of our theories. The interesting part of the story begins when we predicate another well understood *O*-term of the same thing, say, '__ has a negative charge of -1.60×10^{-19} coulombs'. Establishing that the properties, to which these predicates refer, form a fixed relationship attributable to some unobservable *x,* is a considerable epistemic achievement. On the other hand, if this fixed relationship is not epistemically established then the claim, in this regard, is unsubstantiated. This has no bearing on the strength of the claim about what exists, though it does make it a mystery why the relationship is preserved, if indeed it is. Either way, whatever it is that provides the grounding for several mixed predicates such that they form a fixed relation makes structural realism, based on such an approach, look like a strong form realism. The fact that we are claiming of some unobservable *x* that it has a range of properties that form a fixed relation that survives a theory change, and that these properties need not be observable, seems like the kind of claim that any realist might hope to make. There may be a difference between the theoretical commitment to a full blown entity and just those fixed mixed properties we *O*-predicate of a hidden *x;* even so, attributing these properties to a hidden *x* must be too much for structural realism.[17]

Regardless of the implications for structural realism, notice that it is the grounding of the *several* mixed predicates that seems important when considering what is real (rather than considering individually the fixed meaning of each *O*-predicate). And recall that it

[17] Melia and Saatsi [2006] argue that this concession is not too much for structural realism. They maintain that a structural realist can claim that, even if we know of some *x* its velocity, energy and that when placed in a magnetic field it will change direction we still say little about its nature [Melia and Saatsi 2006: 570].

is the project of differential realism to assess what combination of properties form a structure of properties. I see no reason why we cannot form a separate predicate for each property of such a structure, and thus provide the grounding for the fixed relationship between the predicates. We might formulate the relationship between these predicates using functional laws (as outlined in chapter 4, in particular §4.6). According to a differential realist, our assessment of the relationship between several mixed *O*-predicates, as well as determining what we can *O*-predicate of some unobservable, depends on our ability to detect and measure properties.[18] What ever else is required in order to ground the relationship between several fixed mixed predicates so that we might explain their survival through a theory change, it is not clear what the Ramsification of a theory might add. Still, there might be independent advantages to the Ramsey sentence approach. These advantages will be considered in some detail in the next section.

In summing up, the strict structural realists, in so far as they may make use of a Ramsey sentence method, avail themselves of no additional insight by the use of such a method. This is because it is either so weak that we require additional criteria to get the right result, or if strengthened, it is not clear what Ramsification adds to the strict structural realist cause. When considering a stronger version of Ramsification that might support a stronger version of structural realism, it is not clear how 'structural' the realism formed is. This strength requires justification: something the Ramsey approach sentence does not provide. As a result we were left wondering what justified the results of such an approach. I have suggested differential realism helps explain just this support.

6.3 Fixed Reference and Ramsey Sentence Realism

Lewis was right to notice his account requires fixed mixed predicates; unfortunately he provided no substantial analysis of what it is to ground a mixed predicate. The only suggestion he makes is that we must 'understand' *O*-terms and *O*-predicates. Yet a theory might apply well understood mixed predicates regardless of how fixed they are to the entity under analysis. It seems we need some way of grounding these well understood mixed predicates we claim of an unobservable entity, if indeed we can. Lewis just assumes that we have available these fixed mixed *O*-terms and then considers what might follow. Even so, there is something missing from the analysis if there is no demarcation between well understood mixed *O*-predicates and those grounded in a fixed mixed relationship. Differential realism points to precisely this demarcation when analysing what scientists mean when they use theoretical terms. If this is right the situation is somewhat more complex than Lewis envisages. I will argue that it is our ability to ground mixed predicates that is important in certain decisions regarding referential success, and that this is importantly related to empirical success. I will not attempt a full account of reference in science; rather, I hope to show that my

[18] Lewis' appeal to 'old' theory or 'other' theory seems merely to delay the issue.

approach is required of a Ramsey sentence approach that denies the importance of referential success when assessing a theory's approximate truth.[19]

We might begin by considering some of the more usual ways we convey the meaning of theoretical terms. When articulating what the terms of a theory mean, we may use diagrams, graphs, models and familiar analogies. We might also appeal to physical abilities, like the ability to spray something, or an ability to fill a cup with a fluid. Given this enriched view of how we establish the meaning of a term, we might consider what 'caloric' meant for many of those early theorists. These theorists considered caloric a fluid, in part, because this understanding explained a lot of phenomena. Presumably, for these theorists, this had something to do with the analogous and well understood features of more familiar fluids. To this extent the *O*-term 'fluid' has a fixed meaning. But noting this fixed meaning is irrelevant to realism unless it has some relation to what we might detect and measure in virtue of being a fluid. As we will soon see, there is a form of Ramsey sentence realism that argues that it can explain the empirical success of a theory. However, before we consider this, I want to review some of the advantages Lewis claims for his approach. This will help clarify what I am claiming of differential realism and its relation to Ramsey sentence realism.

Lewis claims that his development of the Ramsey sentence, unlike Carnap's, allows for uniqueness in what theoretical terms denote. Lewis [1983] develops his account of uniqueness formally, but the issue at stake can be appreciated by way of example. Early 18th century chemists used the term 'the inflammable gas' as a name (more accurately, a definite description) for a unique kind. However, it so turns out that there are many types of inflammable gas; hydrogen, carbon monoxide, methane and so forth. Now it is true that they are all inflammable, but the term was meant, not just to identify flammability as a property of various gases, but also to identify *the* kind that had a collection of properties that made it unique.[20] Lewis' intuition here would be that, in so far as these early chemists thought they were referring to *one* kind of thing when they used the term 'the inflammable gas', they made a mistake. The fact that several kinds were involved made the uniqueness claim false. Thus the term is denotationless and the theories using this term false.

Of course not all theories imply the uniqueness of the terms involved (consider Faraday's definition of electrical current in the previous chapter). However, Lewis argues that it is a better theory that aspires to uniqueness. That seems defensible, but even with such aspirations, if a theory fails to have a unique referent (in so far as it is multiply realised) do we really have to consider it denotationless? Surely we have more tolerance than this. Unique reference is certainly *preferable* to 'multiple realisation', but

[19] I don't have a fully worked out general theory of reference, though I think Putnam [1970, 1973] is closest to the mark.
[20] This example is somewhat manufactured given that, for many scientists, flammability was a property of phlogiston.

the latter is not a complete failure. I might hope that the lottery numbers *uniquely* match my ticket; indeed, there seems no good reason not to have this preference. However, if my ticket, with several others, matches the winning numbers, even if the winnings are evenly distributed, it still seems that I hold a winning ticket. To accommodate some tolerance here, we might claim it is reasonable that scientists should hope to *eventually* define a theory that does uniquely refer. But when starting out, it is reasonable to achieve 'multiple realisation', for this may be all that is possible. After all, as we start out we may not have the technological ability to detect or measure enough of what is involved such that we can uniquely define the term.

We might see this attitude relevant to the term 'the hydrogen atom'. With the discovery of the isotopes of hydrogen – protium, deuterium and tritium – the term is now known to have multiple realisers. Calling these similar but distinct kinds 'isotopes' must seem rather *ad hoc* under Lewis' schema. If we strictly adhere to the uniqueness requirement, it looks like the term 'the hydrogen atom' does not denote anything. Yet this must seem too conservative for those realists who would consider referential success at least a necessary condition for approximate truth. On the other hand, if we allow some tolerance in our interpretation of what denotes the hydrogen atom, how then to exclude interpreting methane, hydrogen and carbon monoxide as 'isotopes' of the inflammable gas? Lewis does in fact provide a way of adjusting a theory so as to be true or have an interpretation if there is a unique realisation that is *close enough*. However, it is not clear how his advice might adjudicate in these difficult cases.

Differential realism can provide some guidance here. We can start by noting that what the term 'inflammable gas' refers to is well understood. We can accurately measure the flammability of a gas in terms of its flash point relative to the concentration of an oxidising agent. However, it is in virtue of these measurements and their relation to other properties of an inflammable gas that we infer that there are different kinds involved. We might compare this with Faraday's measurement of the (potentially) different electrical fluids. As we saw, Faraday established that there was no measurable difference between the kinds, given a range of specified properties. Because of what his measurements indicated (using the differential approach), it seems reasonable to suppose he has provided us something upon which our theories might further develop. The relationship between the relevant measurable properties established a structure of properties that allows for referential success, even if the full account of the entity or entities involved is not yet available. If, contrary to the facts, a relationship between the flash point, density, solubility, toxicity or whatever was differentially established, it might have been reasonable to regard the term '*the* inflammable gas' as denoting a structure of properties implicit in those theories used in detecting or measuring the relevant properties. If, at some later date, we discovered that somewhat different kinds were involved we might be prepared to treat the new kinds as 'isotopes' of 'the inflammable gas'. It seems, then, that what is doing the work here is the warrant we have in supposing we have a structure of properties, as an initial referent, implicit in the relevant theories. The belief in such a structure is warranted (or not) in virtue of what is indicated by our instruments of detection and measurement using the differential method. As mentioned, we failed to identify the relevant relationship between the measurable or detectable properties of the inflammable gas, and it is this fact that explains why the term 'the inflammable gas' failed to refer.

On the differential realist view, whether our theories correctly or incorrectly refer to what they describe depends on our ability to detect or measure properties. It is risky, though maybe admirable, of a theory to propose a rich structure of properties, and then for scientists to seek to measure what is proposed. Maybe this is where the caloric theorists went wrong. They started out with a rich understanding of the nature of caloric, no doubt informed by a familiar analogy with what it is to be a fluid more generally, but then had to retract features that failed the test of detection or measurement. Had they started out thin and tried to establish what is measurable in the way that Faraday evaluated the measurability of electric current, we might now have a different attitude to the referential success of 'caloric'. Although well understood, it was the fluid-like properties of caloric that, in the end, we failed to detect or measure, and thus why we think the term 'caloric' failed to refer to something fluid-like.[21]

It seems then that there is scope to improve Lewis' programme such that it says more about how to get the right result. David Papineau [1996b] has suggested a plausible way of developing Lewis' approach that might accommodate both the correct reference of theoretical terms, yet allow for principled changes to how we describe them. To capture the required unique referential value, he allows his theoretical term F_i (τ_i in Lewis' account) to be defined with different types of assumptions. Papineau claims that some:

> core assumptions involving F (T_y – "y" for "yes") unquestionably *do* contribute to F's definition, and other accepted assumptions involving F (T_n – "no") unquestionably do *not* contribute, but that beyond that it is indeterminate whether any other generally accepted claims involving F (T_p – "perhaps") have a definitional status [Papineau 1996: 13]. [I have italicised the variables for clarity].

The T_y*'s* guarantee unique reference if there is something unique to refer to, while the various T_p*'s* allow for some development or vagueness of meaning. Obviously the T_n's are those uses of F_i in the postulate that don't contribute anything to F_i's meaning. This division of definitional labour allows for some predicates to be 'criterial' while allowing others to be true or false of what is being referred to. As long as there are enough T_ys, F_i can uniquely refer. The T_ps that refer to non-essential properties allows the total package

[21] Given the range of caloric theories this claim is probably too neat to be true, but it does illustrate how the differential approach might provide grounds for such decisions.

to be vague. This permits us to make discoveries about the kind of thing we are referring to without changing what it is that we are referring to.

Papineau next considers how vague predicates used in science might operate or influence reference under different modal systems. He finds that in general, the use of vague predicates creates no extra problems for the kind of modal claims that scientists make. The principal reason is that the modal claims made by scientists are about natural necessity, not logical or metaphysical necessity (by virtue of convention or rigid designation). Yet, having discovered this, surely we should try and understand the nature of the 'perhaps' according to the modal understanding that scientists do have. Surely there has to be some important relationship between the various *T*'s so that we might understand why a *T* is the kind of *T* it is. It was a T_y that caloric was a fluid, a T_p that it had a mass. This T_p was no accident – fluids as far as we can tell have mass. However, the failure to find this T_p counted against the theory, though for some it didn't. Or consider a different example: once the wave-like nature of light was noticed, something like a range T_y*s* could be given for the luminiferous ether. The development of these T_y*s* was informed by analogous wave-like features noted of familiar materials.[22] Once this was appreciated, a range of T_p's implied by the theory and concepts involved emerged. The ether is either gas-like or *perhaps* like a solid or liquid. *Perhaps,* as a material substance, the velocity of propagation of a wave is relative to the medium.

Papineau's introduction of vague predicates is an improvement, but as it stands it says little about the relation between the various types of assumption. To strengthen the approach Papineau suggests the use of a 'score card' in the assessment of a theory. Here we could score the successful *T's* in terms of their 'definitional importance' – where they fit on the T_y to T_n scale. But how do we make this assessment? By what criterion do we measure definitional importance? Is the final score a comparative measure, or is there some absolute score which ensures reference?

My approach avoids such questions. Some structure between properties has to be proposed using several different, though related, parts of a theory, coupled with the relevant operations of instruments or measurements. Perhaps we can measure or detect the relevant properties according to the differential approach, perhaps not. If we measure or detect the relevant properties we have warrant for the belief that we refer to a structure of properties implicit in the theories involved. However, we may not be able to, at least epistemically, give a full account of the entities involved. We may not know much about these entities other than the measurable properties involved and their structural relation to each other. Analogies and the like can fill out the rest of what we mean, but we need not commit to the full picture. As mentioned before, 'brave' theories can enter the scientific field proposing a rich structure, but they risk a failure to be

[22] Psillos [1999: 130ff.] provides a credible account of development of the development of the theory with the use of these analogies.

detected or measured. According to a differential realist, there has to be a structure of properties to which several *T*'s refer in order that we might judge reference. Pivotal, then, is the establishment of an initial measurable or detectable relationship between properties such that we can have a structure of properties upon which we might build. We can enrich our understanding of what the theoretical term refers to by the addition of more *T*'s (and their relation to the established theory) that range over the same structure. However there is no need for a score card other than the assessment of the differential testing involved. Further structural development might result in 'isotopes' or 'sub-kinds' or a more fine grained structure for the entities involved, though this need not undermine the original referential success. Of course if we fail to measure or detect a structure of properties that several related parts of a theory imply, then the relevant terms of the theory involved fail to refer.

Interestingly, David Papineau and Pierre Cruse have argued that the Ramsey sentence approach allows for a form of scientific realism that does not depend on referential success [Cruse & Papineau 2002]. Cruse [2004] has independently argued for the thesis that the truth of the Ramsey sentence does not rely on reference, but then goes to make a stronger claim regarding irrelevance of reference in any serious issue that a realist should consider. Although these accounts are interesting, I think that they may not entirely remove issues of reference in the way they hope.

According to Papineau and Cruse, since Laudan [1981], realists have understood that successful theories need not refer, so trying to utilise correct reference as an explanation for success is going to be problematic. They then outline various accounts of reference that might get the right result for the realist, but argue that each have their problems. Importantly, unlike the theory itself, a theory's Ramsey sentence does not involve the theoretical terms mentioned in the theory. As we have seen, a Ramsey sentence need only claim that something has the properties described in the theory. According to Papineau and Cruse, a Ramsified theory can hardly fail to refer if it does not *name* the entity that has the properties described in the theory.

> [T]he referential success or failure of the theoretical terms in a theory is irrelevant to the approximate truth of its Ramsey sentence, since those terms do not occur in the Ramsey-sentence [Cruse & Papineau 2002: 178].

The example of a luminiferous aether is then provided to try and clarify the idea of approximate truth. Here the term 'aether' failed to refer, while the Ramsey sentence of the theory using this term may be approximately true, because, '[m]any of the implications of the Ramsey sentence could … still be true' [Cruse & Papineau 2002: 179]. If this is right, the approximate truth of a Ramsey sentence can explain the empirical success of its theory. However, it is not clear how this rather brief suggestion is supposed do what is claimed of it. To help clarify matters, we might provide a plausible account of how such an approach might work.

Cruse and Papineau notice that what is being existentially quantified in a Ramsey sentence is a ‘theoretical term-type’. Using a largely constructed example involving a ‘caloric theory’ we might propose the following.

a. caloric (τ_1) has substance-hood or more specifically fluidity, $F\tau_1$
b. caloric (τ_1) has mass, $M\tau_1$
c. caloric (τ_1) has temperature, $K\tau_1$
d. Adding caloric (τ_1) to an atomic structure (τ_2), expands the atomic structure (given constant pressure), $A\tau_1\tau_2 \rightarrow E\tau_2$

We can formalise the above with the following postulate:

12. $T_c(F\tau_1, M\tau_1, K\tau_1, A\tau_1\tau_2 \rightarrow E\tau_2)$

Ramsifying our postulate we get:

13. $\exists x_1 \exists x_2\, T_c(Fx_1, Mx_1, Kx_1, Ax_1x_2 \rightarrow Ex_2)$

What is being said in (13) is that, given a theory T_c there are the types of things that satisfy the predicates described of them. Anything that satisfies the Ramsey sentence is just the type of thing the theory claims to exist. Yet, if we cannot find anything that satisfies this, what follows? How is it that (13) can be approximately true? Maybe there is another Ramsey sentence that does (we believe) have a satisfier. Say:

13′. $\exists x_1 \exists x_2\, T'_c(\sim Fx_1, \sim Mx_1, Px_1, Kx_1, Ax_1x_2 \rightarrow Ex_2)$

Here (13′) may in fact be true; and, in so far as (13) approximates (13′) we might say it is approximately true. But notice (13′) is not the Ramsey sentence of the theory in question. Moreover, the discovery that (12) is false (or lacks a satisfier) does not make (13′) true or even approximately true. Presumably (13′) is the Ramsey sentence of another theory. Importantly, Cruse and Papineau allow a Ramsey sentence antecedently well understood content. ‘Antecedently understood terms could thus refer to such substantial non-logical relations as causation or correlation, or indeed to many kinds of unobservable things’ [Cruse and Papineau 2002: 182]. Presumably, this sort of content would allow the (approximate?) truth of (13′) to explain the empirical success of its own theory. Thus we can see how certain parts of the old Ramsey sentence are preserved in the new theory, and to this extent we can assess its approximate truth; while we can also explain, using a Ramsey sentence approach, a theory's empirical success.[23]

All this seems plausible. However, one might reasonably ask, what is involved in judging the empirical success that (13′) explains? It seems obvious that a large part of this story will depend upon what we detect and measure, and in virtue of which we attribute properties to some hidden entity that may feature in the explanation. It is true that this need not be the only story: Cruse and Papineau suggest some hypothetical posits are introduced prior to any experiment [Cruse and Papineau 2002: 184]. Even so, empirical success will, in many cases, depend on what we detect and measure. Thus, although the explanatory role of entities goes beyond what we do detect and measure, ultimately something like differential realism may be required to judge whether or not the more hypothetical parts of a theory are empirically assessable. This ability may be important when we come to form the Ramsey sentence (13″) such that we can judge it as explaining the empirical success of its theory, and thus the approximate truth of (13′) and (13). If differential realism does import issues of reference then it's not clear that Cruse and Papineau have avoided all questions of reference. Of course there may be other ways of measuring empirical success without importing issues of reference for this form of Ramsey sentence realism, though Mark Newman [2005] has argued that this is unlikely. So, as it stands Cruse and Papineau's account need not undermine differential realism, and may indeed depend on it.

This issue comes to a head for Cruse [2004], who argues for something stronger; 'there are good reasons to think that questions of realism [and reference] are largely decided by convention and carry no epistemic significance' [Cruse 2004: 133]. To lend weight to this view Cruse provides a thought experiment in which we might change our beliefs about the referential success for particular scientific theories.

> Imagine that instead of deeming the 'aether' a superfluous posit, Albert Einstein had adopted the term to refer to the electromagnetic field, but that Ernest Rutherford had dropped the term 'atom' on the grounds of its connotations of indivisibility in favour of something more descriptive such as 'chemical unit'. Had this happened it is not implausible to suppose that our intuitive judgements about the referential status of 'atom' and 'aether' would have been switched. [Cruse 2004: 137]

Cruse goes on to point out that this switching of our attitude to referential success would not have affected the empirical success of any of the theories involved. Moreover, that

[23] According to Psillos [2009: ch. 9], a Ramsey sentence like (13′), that accounts for the past success (and failure) of previous theories, should be considered as a 'growing existential statement'. It seems this approach is closer to Ramsey's original intention.

this empirical success is assessable using Ramsey sentence approach. If this is right, the referent might seem to fall out of the picture of our assessment of the theory.

Even if we agree with Cruse that there are conventions involved in assessments of referential status, we might suspect that their application should not be entirely arbitrary. As Cruse himself notes, not anything that might satisfy a Ramsey sentence can act as explanation for a theory's empirical success. The fact that some entity on some distant planet might satisfy a Ramsey sentence of a theory that concerns things that happen on this planet seems an inappropriate realiser for what explains the empirical success of the theory. To avoid this problem Cruse offers some suggestive remarks regarding the use of indexical terms such as 'here' and 'this' that, when incorporated into a Ramsey sentence approach, might allow him to avoid this problem. However, it is not clear that this will avoid all issues of reference when assessing certain empirical achievements. This issue is particularly important for those realists who argue for an empirical world that includes the detectable world. According to a differential realist, there are important 'conventions' that get a theory off 'ground level', so to speak. Although it may be true to say that conventions were involved in the evaluation, Faraday was able to test a relationship between the relevant properties so that he might measure an electrical current, and thus he provided us with a way of referring to this kind of unobservable, though detectable, thing. With this minimal structure of properties, implicit in the theories used to measure and detect, the sort of historical revision that Cruse might envisage is unlikely to undermine our intuitions regarding the referential success of the term 'electrical current'. But even if we were to change our mind regarding referential success here, it seems reasonable to suggest that this would threaten what we might count as empirical in this sort of setting. Rather than have Faraday detect and measure an electrical current, a failure to correctly refer here may mean we just have the movement of dials or changes of colour on pieces of paper in a certain experimental setting (or something even less realistic). If issues of reference are not epistemically important it seems this may significantly limit what counts as empirical success. This would be a high price to pay for a form of empirical success without any concern for reference.

Maybe a Ramsey sentence realist can provide a rich account of empirical success that includes the detectable world and does not import issues of reference. However, it seems to me that much work will need to be done here so that they do not beg the question against the anti-realist. Presumably it is a history of important terms failing to refer to what exists in the detectable world that motivates at least some forms of anti-realism. Merely claiming that certain terms are antecedently well understood when determining whatever it is that is detected, and thus what we should count as empirical success, might seem, for some anti-realists, to presuppose the very thing that is at stake.

6.3 Concluding Remarks

When considering the role of the Ramsey sentence approach in developing a viable form of structural realism we found that additional criteria had to be introduced to avoid problems and establish the right result. As a consequence we were left wondering what additional insight the Ramsey sentence approach provided. There are, however, other versions of scientific realism based on a development in the Ramsey sentence approach. But here we found that differential realism helped in areas where their weakness lay. There may be other versions of Ramsey sentence realism that don't require the help of something like differential realism, but I suspect that these will require additional 'machinery' to satisfy the realist. Of course some realists might be happy with the amount of realism provided by an unaided Ramsey sentence approach, but I have argued a scientific realist can have more.

General Conclusions

This thesis has attempted a defence of a form of scientific realism based on what we detect and measure. More specifically, I have defended a form of realism based on what our instruments of detection indicate. Much of the thesis has been developed in contrast to the constructive empiricism offered by Bas van Fraassen. However, it also owes a lot to constructive empiricism; especially in the development of what might count as empirical. Indeed, it is the empirical basis to differential realism that allows it to have some advantages over other forms of realism. Unlike other forms of realism, differential realism does not start with some form of miracle argument regarding the success of science. Rather, it points to the analogous relationship between observation and detection. If the anti-realist were to try and limit the epistemic reach of our instruments of detection and measurement they risks undermining important assumptions in their own account.

Pivotal in the analysis was the development of an informative test. The key intuition behind the test was that, with an independent check on what an instrument indicates we are justified in believing what is indicated. We found that the development of the differential approach by the early operationalists provided important insights into the nature of such a test. Such an approach, however, is not entirely free of a theoretical or conceptual framework. This might have exposed this form of realism to a challenge from the falsificationist who is sceptical of the security of any theory. However, I argued that such scepticism either could not make sense of the entire field of science, or, upon analysis, depended upon a form of secure theory.

One important feature of the thesis is the use of insights provided by the practices of the applied sciences. This approach follows the example of Ian Hacking (and the pragmatists that informed him). However, I have gone further than Hacking. I have gone further because with a secure design theory we can have theoretical knowledge of the entities that populate the detectable and measurable world. Such secure theory does not beg the question against the constructive empiricist, as something analogous is used (according to at least one plausible theory of perception) when forming perceptual judgements about what is observable.

By including what we detect and measure in the applied sciences we were also able to disarm an important anti-realist argument – the pessimistic meta-induction. Once we understood that there is little or no empirical basis for the PMI, we discovered it was more indicative of the differences between the anti-realist and realist. With the PMI dissolved, or at least undermined, a central motive for the development of weaker forms of realism was diminished. There need not be any threat; indeed there might be inductive support for, what our best theories claim to exist in the unobservable realm.

On the realist side of the ledger I have argued that differential realism performs better when considering problems that face other forms of preservative realism (structural realism, deployment realism, semi-realism). It performs better, in large part, because it can give some epistemic guidelines as to how we should interpret what our instruments of detection indicate. These epistemic guidelines would be of limited help to the realist cause, however, without some ontology for us to be realistic about. It is for this reason that I introduced the concept of a structure of properties. Although this term is not used (or not often used) in scientific practice, the concept of a structure of properties seems implicit in what scientists believe is warranted as real as a result of what they detect and measure. This allows those theoretical commitments that account for what is detected and measured to be truth assessable.

Having developed the epistemic guidelines regarding the detectable and measurable world, coupled with a suitable ontology, I argued that recent developments in the Ramsey sentence approach to evaluating theories require something like differential realism. But this just reintroduces issues that the Ramsey sentence approach purports to avoid. I also argue that the structural realist faces serious problems when helping themselves to such an approach.

Although I have had to truncate several issues along the way – leaving my project open to further development – my thesis allows us to have scientific knowledge about the world and the entities within it, even if those entities are unobservable. This allows our scientists, engineers and technicians to have a justification for the beliefs they have about the world. From a philosophical perspective then, my thesis provides epistemic guidelines for judging such beliefs, while providing new insight into the ontology required of such an account.

Bibliography

Achinstein, P. (1968) *Concepts of Science: A Philosophical Analysis*, Johns Hopkins Press, London.

Agassi, J. (1971) *Faraday as a Natural Philosopher*, University of Chicago Press, Chicago & London.

Alspector-Kelly, M. (2004) 'Seeing the Unobservable: van Fraassen and the Limits of Experience', *Synthese*, Vol. 140, pp. 331-353.

Anscombe, G.E.M. (1971) *Causality and Determination*, Cambridge University Press, Cambridge.

Armstrong, D.M. (1997) *A World of States of Affairs*, Cambridge University Press, New York.

Armstrong, D.M. (1973), *Belief, Truth and Knowledge*, Cambridge University Press, London.

Armstrong, D. M. (1978) *Universals and Scientific Realism, Volume II: A Theory of Universals*, Cambridge University Press, Cambridge, New York.

Armstrong, D.M. (1980) 'Against Ostrich Nominalism: A Reply to Michael Devitt', *Pacific Philosophical Quarterly*, Vol. 61, pp. 440-449.

Armstrong, D.M. (1983a) *What is a Law of Nature?* Cambridge University Press, Cambridge.

Armstrong, D.M. (1983b) 'Indeterminism, Proximal Stimuli, and Perception' *The Behavioral and Brain Sciences*, Vol, 6, pp. 55-90.

Armstrong, D.M. (1986) 'In Defence of Structural Universals', *Australasian Journal of Philosophy*, Vol. 64, No. 1, pp. 83-88.

Armstrong, D. (1999) 'A Reply to Ellis', in Sankey, H. (ed.), *Causation and the Laws of Nature*, Kluwer, Dordrecht.

Auxier, R. (1995) 'The Decline of Evolutionary Naturalism in Later Pragmatism', in Hollinger, R. & Depew, D. (eds) *Pragmatism: From Progressivism to Postmodernism*, Praeger, Westport.

Baert, P. (2003) 'Pragmatism, Realism and Hermeneutics', *Foundations of Science*, Vol. 8, pp. 89-106.

Baggott, J. (2004) *Beyond Measure: Modern Physics, Philosophy and the Meaning of Quantum Theory*, Oxford University Press, New York.

Baird, D. (1994) 'Meaning in a Material Medium', *Philosophy of Science Association*, Vol. 2, pp. 441-451.

Baird, D. & Nordman, A. (1994) 'Facts-well-Put', *British Journal for the Philosophy of Science*, Vol. 45, pp. 37-77.

Bigelow, J. and Pargetter, R. (1989) 'A Theory of Structural Universals', *Australasian Journal of Philosophy*, Vol. 67, No.1, pp. 1-11.

Bigelow, J. (1999) 'Scientific Ellisianism', in Sankey, H. (ed.), *Causation and Laws of Nature: Australasian Studies in the History and Philosophy of Science, Vol. 14,* Kluwer Academic Publishing, Dordrecht.

Bence Jones, H. (1870) *The Life and Letters of Faraday*. Longmans, Green & Co. London.

Bohm, D. (2006) *The Special Theory of Relativity*, Routledge Classics, London and New York.

Boyd, R. (1991) 'Realism, Anti-foundationalism and the Enthusiasm for Natural Kinds', *Philosophical Studies*, Vol. 61, pp. 127-148.

Boyd, R. (1984) 'The Current Status of Scientific Realism', in Leplin, J. (ed.) *Scientific Realism*, University of California Press, Berkeley.

Boyd, R. (1985) '*Lex Orandi est Lex Credendi*', in Churchland, P.M. and Hooker, C.A. (eds), *Images of Science; Essays on Realism and Empiricism, with a reply from Bas C. van Fraassen*, University of Chicago Press, Chicago.

Boyd, R. (1990) 'Realism, Approximate Truth, and Philosophical Method', in Savage W.C. (ed.) *Scientific Theories, Minnesota Studies in the Philosophy of Science*, vol. 14, MUP.

Bradley, F.H. (1897) *Appearance and Reality: A Metaphysical Essay*, (second edition) Swan Sonnenschein, London.

Bridgman, P.W. (1938) 'Operational Analysis', *Philosophy of Science*, Vol. 5, No. 2, pp.114-131.

Bridgman, P.W. (1927) *The Logic of Modern Physics*, The Macmillan Co., NY.

Brown, (1985) 'Explaining the Success of Science', *Ratio*, Vol. 27, pp. 49-66.

Buchler, J. (ed.) (1955) *Philosophical Writings of Peirce*, Dover Pub., N.Y.

Bunge, M. (1966) 'Technology as Applied Science', in *Technology and Culture*, Vol. 7, No. 3, pp. 329-347.

Campbell, K. (1990) *Abstract Particulars*, Basil Blackwell, Oxford.

Cao, Y.T. (2003) 'Can We Dissolve Physical Entities into Mathematical Structures?' *Synthese*, Vol. 136, pp. 57-71.

Cao, Y.T. (2003) 'Structural Realism and the Interpretation of Quantum Field Theory', *Synthese*, Vol. 136, pp. 3-24.

Carroll, J.W. (1987) ‘Ontology and the Laws of Nature’, *Australasian Journal of Philosophy*, Vol. 65, No. 3, pp. 261-276.

Carnap, R. (1949) ‘Logical Foundations of the Unity of Science’, in Feigl, H. & Sellars, W. (eds) *Readings in Philosophical Analysis*, Appleton-Century-Crofts, NY.

Carnap, R. (1936-37) ‘Testability and Meaning’, *Philosophy of Science*, Vols. 3 & 4.

Carnap, R. (1966) *Philosophical Foundations of Physics*, Basic Books, New York.

Cartwright N. (1999) *The Dappled World; A Study of the Boundaries of Science*, Cambridge University Press, Cambridge.

Cartwright N. (1980) ‘Do the Laws of Nature State the Facts?’, *Pacific Philosophical Quarterly*, Vol. 61, pp75-84.

Castan□eda, H. (1972) ‘Thinking and the Structure of the World’, *Critica*, Vol. 6, pp. 43-81.

Cavendish, H. (1784) ‘Experiments on Air. By Henry Cavendish, Esq. F. R. S. & S. A.’ *Philosophical Transactions of the Royal Society of London*, Vol. 74, pp. 119-153.

Cei, A. & French, S. (2006) ‘Looking for Structure in All the Wrong Places: Ramsey Sentences, Multiple Realisability, and Structure’, *Studies in the History and Philosophy of Science*. Vol. 37, pp. 633-655.

Chakravartty, A. (1998) ‘Semirealism’, *Studies in History and Philosophy of Science*, Vol. 29, No. 3, pp. 391-408.

Chakravartty, A. (2003) ‘The Structuralist Conception of Objects’, *Philosophy of Science*, Vol. 70, pp.867-878.

Chakravartty, A. (2004) ‘Structuralism as a Form of Scientific Realism’, *International Studies in the Philosophy of Science*, Vol. 18, Nos. 2 & 3, pp. 151-171.

Chang, H. (2003) ‘Preservative Realism and Its Discontents: Revisiting Caloric’, *Philosophy of Science*, Vol. 70, pp. 902-912.

Chalmers, A.F. (2004) *What is Thing Called Science?* (3rd edition), University of Queensland Press, Queensland.

Churchland, P.M. and Hooker, C.A. (eds) (1985), *Images of Science; Essays on Realism and Empiricism, With a Reply From Bas C. van Fraassen*, University of Chicago Press, Chicago.

Churchland, P. (1999) ‘Knowing Qualia: A Reply to Jackson’, in Block, N., Flanagan, O. & Güzeldere, G. (eds) *The Nature of Consciousness: Philosophical Debates*, MIT Press, Cambridge.

Clendinnen, F.J. (1999) ‘Causal Dependence and Laws’, in Sankey, H. (ed.) *Causation and Laws of Nature: Australasian Studies in the History and Philosophy of Science, Vol. 14,* Kluwer Academic Publishing, Dordrecht.

Cohen, S. (1999) ‘Contextualism, Skepticism, and the Structure of Reasons’, *Philosophical Perspectives*, Vol. 13, pp. 57-89.

Collier, J. (1996) ‘On Natural Necessity of Natural Kinds’, in Riggs, P.J. (ed.), *Natural Kinds, Laws of Nature and Scientific Methodology; Australasian Studies in the History and Philosophy of Science, Vol. 12,* pp. 1-10, Kluwer, Dordrecht.

Collier, J. (1988) ‘Supervenience and Reduction in Biological Hierarchies’, in Matthen, M. & Linsky, B. (eds) *Philosophy and Biology; Canadian Journal of Philosophy,* Supplementary Vol. 14, pp.209-234.

Craig, E. (1993) ‘Understanding Scepticism’, in Haldane, J. and Wright, C. (eds), *Reality, Representation and Projection*, Oxford University Press, New York.

Craig, E. (ed.) (1998) *Routledge Encyclopedia of Philosophy*, Routledge, London ; New York.

Cruse, P. (2004) ‘Scientific Realism, Ramsey Sentences and the Reference of Theoretical Terms’, *International Studies in the Philosophy of Science*, Vol. 18, Nos. 2&3, pp. 133-149.

Cruse, P. (2005) ‘Ramsey Sentences, Structural Realism and Trivial Realisation’, *Studies in History and Philosophy of Science*, Vol. 36, pp. 557-576.

Cruse, P. and Papineau, D. (2002) ‘Scientific Realism Without Reference’, in Marsonet, M. (ed.), *The Problem of Realism*, Ashgate, Hampshire.

Curd, M. & Cover, J.A.(eds) (1998) *Philosophy of Science: The Central Issues*, W.W. Norton and Company, New York.

Davis, E.A. and Falconer, I.J. (1997) *J.J. Thomson and the Discovery of the Electron*, Taylor & Francis, London.

Demopoulos, W. & Friedman, M. (1985) ‘Critical Notice: Bertrand Russell’s The Analysis of Matter: Its Historical Context and Contemporary Interest’, *Philosophy of Science*, Vol. 52, pp. 621-639.

Dennett, D.C. (1981) ‘True Believers: The Intentional Strategy and Why it Works’, in Heath, A.F. (ed.) *Scientific Explanation*, Oxford University Press, Oxford.

Dewey, J. (1939) *Logic; The Theory of Inquiry*, Holt, Rinehart & Winston, N.Y.

Dewey, J. (1980) *The Quest For Certainty: A Study of the Relation of Knowledge and Action*, Perigee Print, NY.

Dewey, J. (1981)’Nature, Means and Knowledge’, in Boydston, J.A. (ed.), *John Dewey The Later Works, 1925-1953*, Vol. 1, Southern Illinois University Press, Carbondale and Edwardsville.

Douglas, J.F. (1986) *Solving Problems in Fluid Mechanics: Vol. 1*, Longman Scientific & Technical, Essex.

Doppelt, G. (1973) 'Dretske's Conception of Perception and Knowledge' *Philosophy of Science.* Vol. 40, pp. 433-446.

Dretske, F. (1969) *Seeing and Knowing*, Routledge, London.

Dretske, F. (1977) 'Laws of Nature', *Philosophy of Science,* Vol. 44, pp. 248-268.

Dretske, F. (1981) *Knowledge and the Flow of Information*, MIT Press, Cambridge.

Dretske, F. (1983) 'Précis of *Knowledge and the Flow of Information*', *The Behavioral and Brain Sciences*, Vol, 6, pp. 55-90.

Dretske, F. (1986) 'Misrepresentation', in Bogdan, R.J. (ed.) *Belief: Form, Content and Function*. Clarendon Press, Oxford.

Dretske, F. (1988) *Explaining Behavior: Reasons in a World of Causes*, MIT Press, Massachusetts.

Duhem, P.M. (1954) *The Aim and Structure of Physical Theory,* (Trans.Wiener, P.P.) Princeton U.P., Princeton.

Earman, J. and Roberts, J.T. (2005a) 'Contact with the Nomic: A challenge to the Deniers of Humean Supervenience about Laws of Nature Part I: Humean Supervenience', *Philosophy and Phenomenological Research,* Vol. 71, No. 1, pp. 1-22.

Earman, J. and Roberts, J.T. (2005b) 'Contact with the Nomic: A challenge to the Deniers of Humean Supervenience about Laws of Nature Part I: The Epistemological Argument for Humean Supervenience', *Philosophy and Phenomenological Research,* Vol. 71, No. 2, pp. 253-286.

Eddington, A. (1920) *Space, Time and Gravitation: An Outline of the General Relativity Theory,* Cambridge University Press, Cambridge.

Einstein, A. (1982) 'The Problem of Space, Ether and The Field in Physics', in *Ideas and Opinions,* Crown Pub., NY.

Einstein, A. (2004) *Relativity: The Special and General Theory*, Routledge, London.

Ellis, B. (1987) 'The Ontology of Scientific Realism', in Pettit, P., Sylvan, R. & Norman, J. (eds), *Metaphysics and Morality: Essays in Honour of J.J.C Smart*, Blackwell, Oxford.

Ellis, B. (1996) 'Natural Kinds and Natural Kind Reasoning', Riggs, P.J. (ed.) (1996) *Natural Kinds, Laws of Nature and Scientific Methodology; Australasian Studies in the History and Philosophy of Science, Vol. 12,* Kluwer, Dordrecht.

Ellis, B. (2001) *Scientific Essentialism,* Cambridge University Press, Cambridge.

Ellis, B. & Lierse C. (1994) 'Dispositional Essentialism', *Australasian Journal of Philosophy*, Vol. 72, No.1, pp.27-45, pp. 161-182.

Enc, B (1982) 'Intentional States of Mechanical Devices', *Mind*, Vol. 91, No. 362.

English, J. (1973) 'Underdetermination: Craig and Ramsey', in *Journal of Philosophy*, Vol. 70, pp. 453-463.

Fales, E. (1990) *Causation and Universals*, Routledge, London.

Faraday, M. (1832) 'The Bakerian Lecture – Experimental Researches in Electricity. – Second Series', *Philosophical Transactions of the Royal Society of London*, Vol. 122 (1832) 163 – 194.

Faraday, M. (1832) 'Experimental Researches in Electricity', *Philosophical Transactions of the Royal society of London*, Vol. 122, pp. 125-162.

Faraday, M. (1844) *Experimental Researches in Electricity. Vol II*, Richard and John Edward Taylor, London.

Faraday, M. (1855) *Experimental Researches in Electricity. Vol III*, Richard Taylor & William Francis, London.

Faraday, M. (1920) *Experimental Researches in Electricity*, J.M. Dent and Sons, London.

Feigl, H. & Sellars, W. (eds) (1949) *Readings in Philosophical Analysis*, Appleton-Century-Crofts, New York.

Feyerebend, P. (1970) 'Consolations for the Specialist', in Lakatos, I. & Musgrave, A. (eds) *Criticism and the Growth of Knowledge*, Cambridge University Press, Cambridge.

Feyerabend, P. (1978) *Against Method*, Verso, London.

Fine, A. (1984) 'Einstein's Realism', in Cushing, J.T. Delaney, C.F. and Gutting, G.M. (eds) *Science and Reality: Recent Work in the Philosophy of Science*, UND Press, Notre Dame.

Fine, A. (1984) 'The Natural Ontological Attitude', in Leplin, J. (ed.), *Scientific Realism*, University of California Press, Berkeley.

Fine, A. (2001) 'The scientific Image Twenty Years Later', *Philosophical Studies*, Vol. 106, pp. 107-122.

Fodor, J. (1984) 'Observation Revisited', *Philosophy of Science*, Vol. 51, pp. 23-43.

Foley, R. (1987) 'Dretske's "Information-Theoretic" Account of Knowledge', *Synthese*, Vol. 70, pp. 159-184.

Forge, J. (1986) 'David Armstrong on Functional Laws', *Philosophy of Science*, Vol. 53, pp. 584-7.

Fox. R. (1971) *The Caloric Theory of Gases: From Lavoisier to Renault*, Oxford University Press, London.

French, S. & Ladyman, J. (2003) 'Remodelling Structural Realism: Quantum Physics and the Metaphysics of Structure', *Synthese*, Vol. 136, pp. 31-56.

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

Gillies, D.A. (1972) 'Operationalism', *Synthesis*, Vol. 25, pp. 1-24.

Gillies, D. (1993) 'The Duhem Thesis and the Quine Thesis', *Philosophy of Science in the Twentieth Century, Four Central Themes*, Blackwell, Oxford.

Gillies, D. (1995) 'Popper's Contribution to the Philosophy of Probability' in O'Hear, A. (ed.) *Karl Popper: Philosophy and Problems, Royal Institute of Philosophy Supplement*, Vol. 39, Cambridge University Press, Cambridge.

Glymour, C. (1980) *Theory and Evidence*, Princeton University Press, Princeton.

Glymour, C. (1985) 'Explanation and Realism', in, Churchland, P.M. and Hooker, C.A. (eds) (1985), *Images of Science; Essays on Realism and Empiricism, With a Reply From Bas C. van Fraassen*, University of Chicago Press, Chicago.

Godfrey-Smith, P. (2002) 'Dewey on Naturalism, Realism and Science', *Philosophy of Science,* Vol. 69, No. 3, Supplement: pp. S25-S35.

Godfrey-Smith, P. (2003) *Theory and Reality: An Introduction to the Philosophy of Science*, University of Chicago Press, Chicago.

Gooding, D. (1991) *Science and Philosophy; Experiment and the Making of Meaning*, Kluwer Academic Publishers, Dordrecht.

Gooding, D. & James, F.A.J.L. (eds*)(1985)* *Faraday Rediscovered; Essays on the Life and Work of Michael Faraday*, 1791-1867, Stockton. New York.

Goodman, N (1947) 'The Problem of Counterfactual Conditionals', *The Journal of Philosophy*, Vol. 44, No. 5, pp. 113-128.

Goodman, N. (1954) *Fact Fiction, and Forecast*, Athlone Press, London.

Goodman, N. (1966) *The Structure of Appearance*, Bobbs-Merrill, Indianapolis.

Grice, H.P. (1957) 'Meaning', *Philosophical Review*, Vol. 66, pp. 377-388.

Grünbaum, A (1977) 'Is Psychoanalysis a Pseudo-science: Karl Popper versus Sigmund Freud'**,** *Zeitschrift fuer Philosophische Forschung*, Vol. 31, pp. 333-53.

Hacking, I. (1982), ‘Experimentation and Scientific Realism’, *Philosophical Topics*, Vol. 13, pp. 154-172.

Hacking, I. (1983) *Representing and Intervening: Introductory Topics in the Philosophy of Science*, Cambridge University Press, Cambridge.

Hacking, I. (1985) ‘Do we See Through a Microscope?’, in Churchland, P.M. and Hooker, C.A. (eds) *Images of Science; Essays on Realism and Empiricism, With a Reply From Bas C. van Fraassen*, University of Chicago Press, Chicago.

Hanson, N.R. (1971) ‘On Having the Same Visual Experience’, in Toulmin, S. & Woolf, H. (eds), *What I Don’t Believe and Other Essays*, D. Reidel, Dordrecht.

Hanson, N.R (1972*)* *Patterns of Discovery: An Inquiry into the Conceptual Foundations of Science*, CUP, London.

Hardin, C. and Rosenberg, A. (1982) ‘In Defence of Convergent Realism’, *Philosophy of Science*, Vol. 49, pp. 604-615.

Harré, R. (1981) *Great Scientific Experiments: Twenty Experiments that Changed Our View of the World*, Dover Pub., NY.

Hausman, C.R. (1993) *Charles S. Peirce's Evolutionary Philosophy*, Cambridge University Press, Cambridge.

Hausman, C.R. (2002) ‘Charles Peirce’s Evolutionary Realism as a Process Philosophy’, *Transactions of the Charles S. Peirce Society*, Vol. XXXVIII, No.1-2.

Hempel, C.G. (1952) *Fundamentals of Concept Formation in Empirical Science*, University of Chicago Press, Chicago

Hintikka, J. (1998) ‘Ramsey Sentences and the Meaning of Quantifiers’, *Philosophy of Science*, Vol. 65, pp. 289-305.

Hochberg, H. (1964) ‘Things and Qualities’, in Capitan, W.H. & Merrill, D.D. (eds), *Metaphysics and Explanation*, UPP, Pittsburgh.

Hol, E. (2003) *Empirical Studies on Volatility in International Markets*, Kluwer Academic Pub., Dordrecht.

Hookway, C. (1985) *Peirce*, Routledge & Kegan Paul, London.

Horwich, P. (1991) ‘On the Nature and Norms of Theoretical Commitment’, *Philosophy of Science*, Vol. 58, pp. 1-14.

Howson, C. (ed.) (1976) *Method and Appraisal in the Physical Sciences; The Critical Background to Modern Science, 1800-1905*, Cambridge University Press, Cambridge.

Howson, C. & Urbach, P. (1988) *Scientific Reasoning: The Bayesian Approach*, Open Court, La Salle.

Hume, D. (1992)[1739] *Treatise of Human Nature*, Prometheus Books, New York.

James, W. (1913) *Pragmatism: A New Name For Some Old Ways Of Thinking*, Longman, Green & Co, London.

Kant, I. (1997) *Critique of Pure Reason*, Guyer, P. & Wood, A.W. (eds. trans.), Cambridge University Press, Cambridge.

Ketland, J. (2005) 'Empirical Adequacy and Ramsification', *British Journal for the Philosophy of Science*, Vol 55. pp. 287-300.

Kirkham, R. L. (1998) 'Truth, Pragmatic Theory Of', in Craig, E. (ed.), *Routledge Encyclopedia of Philosophy,* Routledge, London.

Kistler, M. (2006) *Causation and Laws of Nature*, Routledge, London.

Knight, D.M. (1985) 'Davy and Faraday: Fathers and Sons', in Gooding, D. & James, F.A.J.L. (eds*) Faraday Rediscovered; Essays on the Life and Work of Michael Faraday*, 1791-1867, Stockton. NY.

Kornblith, H. (1993) *Inductive Inference and its Natural Ground; An Essay in Naturalistic Epistemology*, MIT Press, Cambridge.

Kuhn, T. (1957) *The Copernican Revolution: Planetary Astronomy in the Development of Western Thought*, Harvard University Press, Cambridge.

Kuhn, T. (1970) *The Structure of Scientific Revolutions (2nd ed)*, University of Chicago Press, Chicago.

Kuhn, T. (1970) 'Logic of Discovery or Psychology of Research?', in Lakatos, I. & Musgrave, A. (eds), *Criticism and the Growth of Knowledge: Proceedings of the International Colloquium in the Philosophy of Science, London 1965, volume 4*, Cambridge University Press, Cambridge.

Kripke, S. (1977) 'Necessity and Identity', in Schwartz, S.P. (ed.) *Naming, Necessity, and Natural Kinds*, Cornell University Press, N.Y.

Kyburg, H.E. (1983) 'Knowledge and the Absolute', *The Behavioral and Brain Sciences*, Vol, 6, pp. 55-90.

Ladyman, J. (1997) 'What is Structural Realism?', *Studies in the History and Philosophy of Science*, Vol. 29, No. 3, pp. 409-424.

Ladyman, J. (2000) 'What is Really Wrong With Constructive Empiricism? Van Fraassen and the Metaphysics of Modality', *British Journal for the Philosophy of Science*, Vol. 5.

Lakatos, I. (1978) *The Methodology of Scientific Research Programmes: Philosophical Papers Volume 1*, Worrall, J. & Currie, G. (eds), Cambridge University Press, London.

Lakatos, I. (1999) 'Lecture Four: 'Comparing Demarcation Criteria: Verificationism and Conventionalism', in Motterlini, M. (ed.), *For and Against Method: Including Lakatos's Lectures of Scientific Method and the Lakatos – Feyerabend Correspondence*, University of Chicago Press, Chicago.

Lakatos, I. & Musgrave, A. (1970), *Criticism and the Growth of Knowledge: Proceedings of the International Colloquium in the Philosophy of Science, London 1965, Volume 4*, Cambridge University Press Cambridge.

Lange, M. (1992) 'Armstrong and Dretske on the Explanatory Power of Regularities' *Analysis*, Vol. 52, No. 3, pp. 154-159.

Lange, M (2000) *Natural Laws in Scientific Practice*, Oxford University Press, New York.

Lange, M. (2002) 'Baseball, Pessimistic Inductions and the Turnover Fallacy', *Analysis*, Vol. 62, No.4, pp. 282-285.

Lange, M. (2004) 'A Note on Scientific Essentialism, Laws of Nature, and Counterfactual Conditionals', *Australasian Journal of Philosophy*, Vol. 82, No. 2, pp. 227-241.

Lange, M. (2005a) 'Reply to Ellis and to Handfield on Essentialism, Laws, and Counterfactuals', *Australasian Journal of Philosophy*, Vol. 83, No. 4, pp. 581-588.

Lange, M. (2005b) 'Laws and Their Stability', *Synthese*, Vol. 144, pp.415-432.

Laudan, L. (1981) 'A Confutation of Convergent Realism', *Philosophy of Science*, Vol 48, pp. 19-49.

Laudan, L, (1983) 'The Demise of the Demarcation Problem' in Cohen, R. and Laudan, L.(eds), *Physics, Philosophy, and Psychoanalysis: Essays in Honor of Adolf Grünbaum*, D. Reidel Pub. Co., Dordrecht.

Laudan, L. (1984) 'Realism Without the Real', *Philosophy of Science*, Vol. 51, No.1, pp. 156-162.

Laudan, L. (1990) 'Demystifying Underdetermination', in Savage, W.C. (ed.) *Scientific Theories*, vol. 14, *Minnesota Studies in the Philosophy of Science*, UMP, Minneapolis.

Laudan, L. (1996) 'Progress or Rationality? The Prospects for Normative Naturalism' in Papineau, D. (ed.), *The Philosophy of Science*, OUP, Oxford.

Laymon, R. (1984) 'The Path From to Theory', in Leplin, J. (ed.) *Scientific Realism*, University of California Press. Berkeley.

Legg, C. (1999) 'Real Law in Charles Peirce's "Pragmaticism"', in Sankey, H. (ed.), *Causation and Laws of Nature: Australasian Studies in the History and Philosophy of Science, Vol. 14,* Kluwer Academic Publishing, Dordrecht.

Leplin, J. (1987) 'Surrealism', *Mind*, Vol. 96, pp. 519-524.

Lewis, D. (1972) 'Psychophysical and Theoretical Identifications', *Australasian Journal of Philosophy,* Vol. 3, pp. 249-258.

Lewis, D. (1973) *Counterfactuals*, Blackwell, Oxford.

Lewis, D. (1983) 'How To Define Theoretical Terms', from his *Philosophical Papers; Volume 1*, Oxford University Press, N.Y.

Lewis, D. (1983b) 'New Work for a Theory of Universals', *Australasian Journal of Philosophy,* Vol. 61, pp. 343-377.

Lewis, D. (1986) 'Against Structural Universals', *Australasian Journal of Philosophy,* vol. 64, pp. 25-46.

Lewis, D. (1994) 'Humean Supervenience Debugged', *Mind,* Vol. 103, pp. 473-490.

Lewis, D. (1999) 'Finkish Dispositions', in Sankey, H. (ed.) *Causation and Laws of Nature: Australasian Studies in the History and Philosophy of Science, Vol. 14,* Kluwer Academic Publishing, Dordrecht.

Lewis, P. (2001) 'Why the Pessimistic Induction is a Fallacy', *Synthese,* Vol. 129, pp. 371-380.

Lindsay, R.B. (1937) 'A Critique of Operationalism in Physics', *Philosophy of Science*, Vol. 4.

Lyons, T.D. (2002) 'Scientific Realism and the Pessimistic Meta-Modus Tollens', in Clark, S. & Lyons, T.D. (eds) *Recent Themes in the Philosophy of Science: Scientific Realism and Commonsense*, Kluwer Academic Pub., Dordrecht.

Maher, P. (1990) 'Acceptance Without Belief', *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association, Vol. 1, pp. 381-392.*

Martin, C.B. (1994) 'Dispositions and Conditionals', *The Philosophical Quarterly,* Vol. 44. No. 174, pp. 1-8.

Maund, B. (2003) *Perception,* Acumen, Chesham.

Maxwell G. (1962) 'The Ontological Status of Theoretical Entities', in Feigl, H. & Maxwell, G., (eds), *Scientific Explanation, Space and Time, Vol 3. Minnesota Studies in the Philosophy of Science,* University of Minneapolis Press, Minneapolis.

Maxwell, G. (1970a) 'Theories Perception and Structural Realism', in Colodny, R.G. (et al) (eds) *The Nature & Function of Scientific Theories : Essays in*

Contemporary Science and Philosophy, University of Pittsburgh Press, Pittsburgh.

Maxwell, G. (1970b) "Structural Realism and the Meaning of Theoretical Terms', in Winokur, S. & Radner, M. (eds) *Minnesota Studies in the Philosophy of Science, Vol. 4,* University of Minnesota Press, Minneapolis.

Maxwell, J.C. (1965a) 'On Faraday's Lines of Force', in Niven, W.D. (ed.) *The Scientific Papers of James Clerk Maxwell,* Dover, N.Y.

Maxwell, J.C. (1965b) 'A Dynamical Theory of Electromagnetic Field', in Niven, W.D. (ed.) *The Scientific Papers of James Clerk Maxwell*, Dover, N.Y.

Maxwell, N. (1972) 'A Critique of Popper's Views on Scientific Method', *Philosophy of Science,* June 1972, Vol. 39, pp. 131-153.

Mayo, D. (1991) 'Novel Evidence and Severe Tests', *Philosophy of Science,* Vol. 58, pp. 523-552

Mayo, D. (1994) 'The New Experimentalism, Topical Hypotheses, and Learning from Error', in, *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association,* Vol. 1, pp. 270-279.

Mayo, D. (1996a) 'Ducks, Rabbits, and Normal Science: Recasting the Kuhn's-eye View of Popper's Demarcation of Science', *British Journal for the Philosophy of Science*, Vol. 47, pp. 271-291.

Mayo, D. (1996b) *Error and the Growth of Experimental Knowledge,* University of Chicago Press, Chicago.

Mayo, D. (1997) 'Duhem's Problem, the Bayesian Way, and Error Statistics, or "What's Belief Got to Do With It?"', *Philosophy of Science,* Vol. 64, pp. 222-244.

Melia, J. & Saatsi, J. (2006) 'Ramseyfication and Theoretical Content', *British Journal for the Philosophy of Science,* Vol. 57, pp. 561-585.

Mellor, D.H. (1980) 'Necessity and Universals in Natural Laws', in Mellor, D.H. (ed.) *Science, Belief and Behaviour: Essays in Honour of R.B. Braithwaite,* Cambridge University Press, Cambridge.

Mellor, D.H. (2000) 'The Semantics and Ontology of Dispositions', *Mind,* Vol. 109, pp. 757-780.

Mellor, D.H. & Oliver, A. (eds) (1997) *Properties*, Oxford University Press, Oxford.

Menuge, A. (1995) 'The Scope of Observation', *The Philosophical Quarterly,* Vol. 45, No. 178, pp.60-69.

Miggotti, M. (1998) 'Peirce's Double Aspect theory of Truth', in Misak, C.J. (ed.), *Pragmatism*: *Canadian Journal of Philosophy,* Supplementary Vol 24.

Milgrom, M. (2002) 'Does Dark Matter Really Exist?' *Scientific American*, Aug 2002, Vol. 287 Issue 2.

Monton, B. and van Fraassen, B. (2003) 'Constructive Empiricism and Modal Nominalism', *British Journal for the Philosophy of Science*, Vol. 54, pp.405-422.

Morganti, M. (2004) 'On the Preferability of Epistemic Structural Realism', *Synthese*, Vol. 142, pp. 81-107.

Musgrave, A. (1975) 'Popper and "Diminishing Returns from Repeated Tests"', *Australasian Journal of Philosophy*, Vol. 53, pp. 248-53.

Musgrave, A. (1976) 'Why did Oxygen Supplant Phlogiston? Research Programmes in the Chemical Revolution', in Howson, C. (ed.) *Method and Appraisal in the Physical Sciences; The Critical Background to Modern Science, 1800-1905,* Cambridge University Press, Cambridge.

Musgrave, A. (1985) 'Realism Versus Constructive Empiricism', in Churchland, P.M. and Hooker, C.A. (eds) (1985), *Images of Science; Essays on Realism and Empiricism, With a Reply From Bas C. van Fraassen*, University of Chicago Press, Chicago.

Musgrave, A. (1988) 'The Ultimate Argument for Scientific Realism', in Nola, R. (ed), *Relativism and Realism in Science,* Kluwer, Dordrecht.

Neresessian, N.J. (1985) 'Faraday's Field Concept', in Gooding, D. & James, F.A.J.L.(eds*) Faraday Rediscovered; Essays on the Life and Work of Michael Faraday*, 1791-1867, Stockton. New York.

Newton-Smith, W.H. (1978) 'The Underdetermination of Theory by Data', *Proceedings of the Aristotelian Society*, Supp. Vol. 52, pp. 71-91.

Newton-Smith, W.H. (1981) *The Rationality of Science*, Routledge & Kegan Paul, Boston.

Newman, M. (2005) 'Ramsey Sentence Realism as an Answer to the Pessimistic Meta-Induction', *Philosophy of Science*, Vol. 75, No. 5, pp. 1373-1384.

Newman, M.H.A. (1928) 'Mr. Russell's "Causal Theory of Perception"', *Mind*, Vol. 37, No. 146. pp. 137-148.

Niven, W.D. (ed.) (1965), *The Scientific Papers of James Clerk Maxwell*, Dover, N.Y.

Noë, A. (2004) *Action in Perception,* MIT Press, Cambridge Mass.

O'Hear, A. (1975) 'Rationality of Action and Theory Testing in Popper', *Mind*, Vol. 84, No. 334, pp. 273-6.

Øhrstrøm, P. and Hasle, P.F.V. (1995) *Tense Logic: From Ancient Ideas to Artificial Intelligence*, Kluwer, Dordrecht.

Pagès, J. (2002) 'The Dretske-Tooley-Armstrong Theory of Natural Laws and the Inference Problem', *International Studies in the Philosophy of Science,* Vol. 16, No. 3, pp. 227-243

Papineau, D. (ed.) (1996a) *The Philosophy of Science,* Oxford University Press,

Papineau, D. (1996b) 'Theory-dependent Terms', *Philosophy of Science,* Vol. 63, pp. 1-20.

Partington, J.R. (1962) *A History of Chemistry, Volume Three,* Macmillan & Co, London.

Peirce, C.S. (1955) 'Synechism, Fallibilism, and Evolution', in Buchler, J. (ed.) *Philosophical Writings of Peirce,* Dover Pub., N.Y.

Peirce, C.S. (1960), "Some Consequences of Four Incapacities", in Hartshorne, C. and Weis,P. (eds) *Collected Papers of Charles Saunders Peirce,* Vol. V, Harvard University Press, Cambridge Mass.

Peirce, C.S. (1998), 'The Fixation of Belief', in Cohen, M.(ed.) *Chance Love and Logic: Philosophical Essays,* UNP, Lincoln Nebraska.

Poincaré, H. (1952) *Science and Hypothesis*, Dover Publications, New York.

Popper, K. (1963) *Conjectures and Refutations; The Growth of Scientific Knowledge*, Routledge, London.

Popper, K. (1965) *The Logic of Scientific Discovery,* Harper & Row, NY.

Popper, K. (1972) *Objective Knowledge, An Evolutionary Approach,* Clarendon Press, Oxford.

Popper, K (1974) "Replies To My Critics", in Schilpp, P.A. (ed.) *The Philosophy of Karl Popper,* pp. 961-1197, Open court, La Salle.

Preston, J. (1994) 'Methodology, Epistemology and Conventions: Popper's Bad Start', in *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association,* Vol. 1, pp. 314-322.

Priestley, J. (1775) *Experiments and Observations on Different Kinds of Air, Vol. 2,* Printed for J. Johnson, London.

Priestley, J. (1901) 'On Dephlogisticated Air', in *The Discovery of Oxygen, Vol. 1,* Alembic Club, Edinburgh.

Priestley, J. (1970) *Heads of Lectures on a Course of Experimental Philosophy, Particularly Including Chemistry,* Kraus Reprint, NY.

Psillos, S. (1995) 'Is Structural Realism the Best of Both Possible Worlds?', *Dialectica,* Vol. 49, No. 1, pp. 15-46.

Psillos, S. (1996) 'On van Fraassen's Critique of Abductive Reasoning', *The Philosophical Quarterly*, Vol. 46, pp. 31-47.

Psillos, S. (1999), *Scientific Realism: How Science Tracks the Truth*, Routledge, London and N.Y.

Psillos, S. (2000) 'Rudolf Carnap's 'Theoretical Concepts in Science'', *Studies in History and Philosophy of Science*, Vol. 32, No.1, pp. 151-172.

Psillos, S. (2001) 'Is Structural Realism Possible?', *Philosophy of Science,* Vol. 68, pp. 13-24.

Psillos, S. (2006) '*The* Structure, the *Whole* Structure, and Nothing *but* the Structure?' *Philosophy of science*, Vol. 73, pp. 560-570.

Psillos, S. (2009) *Knowing the Structure of Nature: Essays on Realism and Explanation*, Palgrave Macmillan, New York.

Putnam, H. (1970) 'Is Semantics Possible', in Keifer, H.E. & Munitz, M.K. (eds) *Language Belief and Metaphysics,* State University of New York Press, Albany.

Putnam, H. (1973) 'Meaning and Reference', *The Journal of Philosophy*, Vol. LXX, pp. 699-711.

Putnam, H. (1975) 'The meaning of 'meaning'', from his *Mind, Language and Reality: Philosophical Papers, Vol. 2,* Cambridge University Press, Cambridge.

Putnam, H. (1978) *Meaning and the Moral Sciences*, Routledge & Kegan Paul, London.

Quay, P.M. (1974) 'Progress as a Demarcation Criterion for the Sciences', *Philosophy of Science*, Vol. 41, pp.154-170.

Quine, W.V. (1969) *Ontological Relativity and Other Essays*, Columbia University Press, NY.

Ramsey, F.P. (1929/1990) 'Theories', in Mellor, D.H. (ed.) *Philosophical Papers: F.P. Ramsey,* Cambridge University Press, Cambridge.

Redhead, M. (2001) 'Quests of a Realist', *Metascience*, Vol. 10, No. 3, pp. 341-347.

Resnik, D.B. (1994) 'Hacking's Experimental Realism', *Canadian Journal of Philosophy*, Vol. 24, No. 3, pp. 395-412.

Resnik, D. (2000) 'A Pragmatic Approach to the Demarcation Problem', *Studies in History and Philosophy of Science*, Vol. 31, No. 2, pp. 249-267

Riggs, P.J. (ed.) (1996) *Natural Kinds, Laws of Nature and Scientific Methodology; Australasian Studies in the History and Philosophy of Science, Vol. 12,* Kluwer, Dordrecht.

Rorty, R. (1982) *Consequences of Pragmatism (Essays: 1972-1980),* UMP, Minneapolis.

Rothbart, D. (1998) 'Extending Popper's Epistemology to the Lab', *Dialectica*, Vol. 52, No. 3.

Russell, B. (1927) *The Analysis of Matter,* Allen & Unwin, London.

Russell, B. (1948) *Human Knowledge, Its Scope and Limits*, Simon and Schuster, New York.

Saatsi, J.T. (2005) 'On the Pessimistic Induction and Two Fallacies', *Philosophy of Science,* Vol. 72, pp. 1088–1098.

Salmon, W.C. (1981) 'Rational Prediction', *British Journal for the Philosophy of Science,* Vol. 32. pp. 115-125.

Sankey, H. (ed.) (1999), *Causation and Laws of Nature: Australasian Studies in the History and Philosophy of Science, Vol. 14,* Kluwer Academic Publishing, Dordrecht.

Scheffler, I. (1974) *Four Pragmatists: A Critical Introduction to Peirce, James, Mead and Dewey*, Humanities Press, New York.

Schlesinger, G. (1958) 'P.W. Bridgman's Operational Analysis: The Differential Aspect' *British Journal for the Philosophy of Science,* Vol. 9, pp. 299-306.

Smart, J.J.C. (1963) *Philosophy and Scientific Realism,* Routledge & Kegan Paul, New York.

Smart, J.J.C. (1993) 'Laws of Nature as a Species of Regularities', in Bacon, J., Campbell, K. and Reinhardt, L. (eds), *Ontology, Causality and Mind: Essays in Honour of D.M. Armstrong,* Cambridge University Press, Cambridge.

Sober, E. (1982) 'Dispositions and Subjunctive Conditionals, or, Dormative Virtues Are No Laughing Matter', *The Philosophical Review,* Vol. XCI, No. 4, pp. 591- 596.

Stampe, D. (1977) 'Toward a Causal theory of Linguistic Representation', in French, O., Uehling, T. & Wettstein, H. (eds) *Midwest Studies in Philosophy*, Vol. 2, University of Minnesota Press.

Stanford, P.K. (2003) 'No Refuge for Realism: Selective Confirmation and the History of Science', *Philosophy of Science,* Vo. 70, pp. 913-925.

Suppe, F. (1972) 'Theories, Their Formulations, and the Operational Imperative, *Synthese,* Vol. 25, pp. 129-164.

Teller, P. (1995) *An Interpretive Introduction to Quantum Field Theory***,** Princeton University Press, Princeton.

Teller, P. (2001) 'Whither Constructive Empiricism?', *Philosophical Studies,* Vol. 106, pp. 123-150.

Tristan, J. (2002) 'The Role of Measurement in Inquiry', in Burke, T.F., Hester, D.M. & Talisse, R.B. (eds) *Dewey's Logical Theory: New Studies and Interpretations*, Vanderbilt University Press, Nashville.

Tooley, M. (1977) 'The Nature of Laws', *Canadian Journal of Philosophy*, Vol. 7, pp. 667-698.

Tooley, M. (1987) *Causation: A Realist Approach*, Clarendon, Oxford.

Tweney, R.D. (1985) 'Faraday's Discovery of Induction: A Cognitive Approach', in Gooding, D. & James, F.A.J.L., (eds) *Faraday Rediscovered; Essays on the life and Work of Michael Faraday*, 1791-1867, Stockton. NY.

Vaihinger, H. (1968) *The philosophy of 'As if' : a System of the Theoretical, Practical and Religious Fictions of Mankind*, Ogden, C.K. (trans.) Routledge & Kegan Paul, London.

Van Cleve, J (1998) 'Three Versions of Bundle Theory', in Laurence, S. & Macdonald, C. (eds), *Contemporary Readings in the Foundations of Metaphysics*, Blackwell, Oxford.

Van Fraassen, B.C. (1980) *The Scientific Image*, Clarendon Press, Oxford.

Van Fraassen, B.C. (1981) 'Essences and the Laws of Nature', in Healey, R. (ed.) *Reduction, Time and Reality: Studies in the Philosophy of the Natural Sciences*. CUP, Cambridge.

Van Fraassen, B. (1985) 'Empiricism in the Philosophy of Science', in Churchland, P.M. and Hooker, C.A. (eds) (1985) *Images of Science; Essays on Realism and Empiricism, with a Reply from Bas C. van Fraassen*, University of Chicago Press, Chicago.

Van Fraassen, B.C. (1987) 'Armstrong on Laws and Probabilities', *Australasian Journal of Philosophy,* Vol. 65, No. 3, pp. 243-260.

Van Fraassen, B.C (1989) *Laws and Symmetry*, Oxford University Press, New York.

Van Fraassen, B. (1992) 'From Vicious Circle to Infinite Regress, and Back Again', *Philosophy of Science Association*, Vol. 2, pp. 6-29.

Van Fraassen, B. (1994) 'Against Transcendental Empiricism', in Stapelton, T.J. (ed.) *The Question of Hermeneutics: Essays in Honour of Joseph J. Kockelmans*, Kluwer Academic Publishers, Dordrecht.

Van Fraassen, B. (2006) 'Structure: Its Shadow and Substance', *British Journal for the Philosophy of Science*, Vol. 57, pp. 275-307.

Van Fraassen, B. (2008) *Scientific Representation: Paradoxes of Perspective*, Oxford Scholarship Online.

Votsis, I. (2003) 'Is Structure Not Enough?', *Philosophy of Science,* Vol. 70, pp. 879-890.

Wheeler, G.R. (2000) 'Error Statistics and Duhem's Problem', *Philosophy of Science,* Vol. 67, pp. 410-420.

Whitaker, A. (1996) *Einstein, Bohr and the Quantum Dilemma,* Cambridge University Press, Cambridge.

Whittaker, E. (1951) *A History of the Theories of Aether and Electricity,* Thomas Nelson, London.

Williams L.P. (1965) *Michael Faraday,* Chapman and Hall, London.

Williams L.P. (ed.) (1971) *The Selected Correspondence of Michael Faraday,* Vol.1 1812-1848, Cambridge University Press, Cambridge.

Williams, L.P. (1985) 'Faraday and Ampère: A Critical Dialogue', in Gooding, D. & James, F.A.J.L. (eds*) Faraday Rediscovered; Essays on the Life and Work of Michael Faraday*, 1791-1867, Stockton. NY.

Wilson, M. (1985) 'What Can Theory Tell Us about Observation?', in Churchland, P.M. and Hooker, C.A. (eds) (1985), *Images of Science; Essays on Realism and Empiricism, With a Reply From Bas C. van Fraassen,* University of Chicago Press, Chicago.

Winnie, J.A. (1967) 'The Implicit Definition of Theoretical Terms', *British Journal for the Philosophy of Science,* Vol. 18, pp. 223-229.

Worrall, J. (1989) 'Structural Realism the Best of Both Worlds', *Dialectica,* Vol. 43, pp. 99-124.

Worrall, J. (2007) 'Miracles and Models: Why Reports of the Death of Structural Realism May Be Exaggerated', in O'Hear, A. (ed.) *Philosophy of Science: Royal Institute of Philosophy Supplement,* Vol. 61, Cambridge University Press, Cambridge.

Zahar, E. (1976) 'Why did Einstein's Programme Supersede Lorentz's?', in Howson, C. (ed.) (1976) *Method and Appraisal in the Physical Sciences; The Critical Background to Modern Science, 1800-1905,* Cambridge University Press, Cambridge.

Zahar, E. & Worrall, J. (2001) 'Ramsification and Structural Realism', in Zahar, E. *Poincaré's Philosophy: From Conventionalism to Phenomenology,* Open Court, Chicago.

End