

**Examining Scientific Perspectivism:
Instruments, Models, and Kinds**

Submitted by Kane Finlay Baker to the University of Exeter
as a thesis for the degree of Doctor of Philosophy in Philosophy,
in March 2022.

This thesis is available for library use on the understanding that it is copyright material and that no quotation from the thesis may be published without proper acknowledgement.

I certify that all material in this thesis which is not my own work has been identified and that any material that has previously been submitted and approved for the award of a degree by this or any other University has been acknowledged.

Abstract

This thesis is an examination of perspectivism in philosophy of science, an approach proposed as a *via media* between standard realist and antirealist views. The thesis is structured into four parts. Part I introduces the varieties of perspectivism in the literature and outlines the specific form of perspectivism that is the focus of this thesis, which is based on the work of Ronald Giere. I distinguish Giere's perspectivism from traditional realism and outline the central challenge to the position, a challenge that I call "Escape from Perspective". The next two parts are critical. In Part II, I examine Giere's perspectivist account of instruments ("instrumental perspectivism"). I find that Giere's arguments for instrumental perspectivism fail, and that instrumental perspectivism is dependent on a perspectivist account of theories and models ("theoretical perspectivism") in a way that robs it of philosophical significance. If this is correct, the question is simply whether theoretical perspectivism can be defended. This question is taken up in Part III. This part focuses on two arguments for perspectivism, the incompatible models argument and the failure of fit argument. I argue that the incompatible models argument is unconvincing, but that the failure of fit argument presents a significant challenge to standard realist views. The rest of Part III attacks the potential realist responses to this argument. Finally, in Part IV, I turn to the constructive project of developing a perspectivist approach that respects the insights of the failure of fit argument. In this part, I present what I take to be the most defensible form of perspectivism, and I argue that it has the resources to meet the Escape from Perspective challenge. Although my perspectivism is different from Giere's in many respects, I end by showing that it exhibits all the characteristics distinctive of perspectivism, as Giere defined it.

Contents

| | |
|---|-----------|
| I: INTRODUCTION | 7 |
| Chapter 1: Perspectivism | 7 |
| 1.1. Perspectivisms..... | 7 |
| 1.2. Giere's perspectivism..... | 9 |
| 1.3. The colour analogy..... | 13 |
| 1.4. Six features of colour perspectives..... | 16 |
| 1.5. Escape from Perspective..... | 22 |
| 1.6. Historical parallels..... | 24 |
| 1.7. Outline of the thesis..... | 29 |
| II: INSTRUMENTAL PERSPECTIVISM | 31 |
| Chapter 2: Partiality and Opacity | 31 |
| 2.1. Introduction..... | 31 |
| 2.2. Instruments as causal processes..... | 33 |
| 2.2.1. Varieties of instruments..... | 34 |
| 2.2.2. Nomological machines..... | 36 |
| 2.3. The direct argument: Partiality and opacity..... | 39 |
| 2.4. Do instruments distort?..... | 44 |
| 2.5. Three puzzles about instrumental perspectives..... | 53 |
| 2.5.1. The proliferation of perspectives..... | 53 |
| 2.5.2. The indeterminacy of perspectives..... | 55 |
| 2.5.3. Interacting perspectives..... | 58 |
| Chapter 3: Extended Perception | 61 |
| 3.1. Do instruments extend perception?..... | 61 |
| 3.2. An alternative view of instruments..... | 64 |
| 3.3. Why the appeal to the parity principle fails..... | 69 |
| 3.4. Beyond the parity principle?..... | 74 |
| 3.5. Conclusion..... | 75 |

| | |
|---|------------|
| III: THEORETICAL PERSPECTIVISM..... | 77 |
| Chapter 4: Incompatible Models and Perspectival Truth..... | 77 |
| 4.1. Multiplicity..... | 77 |
| 4.2. The argument from incompatible models..... | 78 |
| 4.3. Moderate perspectivism..... | 84 |
| 4.3.1. Similarity..... | 84 |
| 4.3.2. Truth..... | 94 |
| 4.4. Relativist perspectivism..... | 102 |
| 4.5. The split strategy..... | 108 |
| 4.6. The coherence of relativism..... | 113 |
| 4.7. Alternative accounts of incompatible models..... | 117 |
| Chapter 5: Failure of Fit..... | 125 |
| 5.1. Language and reality..... | 125 |
| 5.2. Taking a perspective on Putnam's world..... | 131 |
| 5.3. Complexity..... | 136 |
| 5.4. Incompatible models again..... | 143 |
| Chapter 6: The Language of the World..... | 146 |
| 6.1. Introduction..... | 146 |
| 6.2. Denial..... | 146 |
| 6.3. Approximation..... | 150 |
| 6.4. Evolution..... | 151 |
| 6.5. Propositions..... | 155 |
| 6.6. Mathematization..... | 156 |
| 6.7. Natural kinds..... | 163 |
| Chapter 7: Natural Kinds..... | 165 |
| 7.1. Introduction..... | 165 |
| 7.2. Traditional approaches to classification..... | 166 |
| 7.3. Perspectivist accounts of kinds..... | 175 |
| 7.4. Cretu's perspectivist account..... | 179 |

| | |
|---|------------|
| IV: PERSPECTIVISM AND KINDS..... | 188 |
| | |
| Chapter 8: Theoretical Kinds..... | 188 |
| 8.1. Introduction..... | 188 |
| 8.2. Giere on theoretical kinds..... | 188 |
| 8.3. Idealizations..... | 194 |
| 8.4. Schematic objects..... | 195 |
| 8.5. Boundary construction..... | 203 |
| 8.5.1. Chunking continuous variation..... | 203 |
| 8.5.2. Causal/constitution indeterminacy..... | 205 |
| 8.5.3. Scale selection..... | 207 |
| 8.6. Coordination..... | 211 |
| 8.7. Social constructivism..... | 220 |
| | |
| Chapter 9: Escape from Perspective..... | 223 |
| 9.1. Introduction..... | 223 |
| 9.2. Escape from Perspective..... | 224 |
| 9.2.1. Distortion again..... | 225 |
| 9.2.2. Dispositions..... | 228 |
| 9.2.3. Behind-the-schemes..... | 234 |
| 9.3. A general objection..... | 239 |
| 9.4. Incoherence?..... | 245 |
| | |
| Chapter 10: Towards a New Perspectivism..... | 250 |
| 10.1. Introduction..... | 250 |
| 10.2. Interactionism..... | 252 |
| 10.3. Truth-value..... | 252 |
| 10.3.1. Correspondence theories..... | 253 |
| 10.3.2. Deflationist theories..... | 255 |
| 10.3.3. Coherence theories..... | 256 |
| 10.3.4. Pragmatist theories..... | 257 |
| 10.4. Perspectives track regularities in the world..... | 258 |
| 10.5. Perspectives are partial..... | 260 |
| 10.6. Different perspectives may both be accurate..... | 260 |

10.7. Perspectives and purposes..... 261

REFERENCES..... 263

I

INTRODUCTION

Chapter 1

Perspectivism

1.1. Perspectivisms

This thesis is an examination of scientific perspectivism. Perspectivism was introduced into recent philosophy of science by Ron Giere (1999, 2006), who conceived of it as a kind of limited realism, a *via media* between what he called “objectivist realism” and various forms of antirealism such as social constructivism. The basic idea of perspectivism is that our best scientific theories deliver reliable knowledge of the world, but that this is always relative to scientific perspectives, where multiple perspectives are in principle always available on a given domain. In particular, perspectivism emphasizes what we might call the “multiplicity intuition”, that the truth about a domain is revealed from multiple different viewpoints.

It is worth stating at the outset that an immediate problem with any examination of perspectivism is that this basic idea has been developed in a variety of ways. As a result, it’s not clear what perspectivism even is. Various approaches have been proposed in the recent literature (see for example the essays in Massimi and McCoy 2019). Different philosophers, both those working under the “perspectivist” label and those criticizing the view, will give different answers to foundational questions: How does perspectivism relate to realism, pluralism, and pragmatism? What is the perspectivist account of truth? Is perspectivism committed to epistemic relativism? Is perspectivism committed to any particular position with respect to the structure of scientific theories (the syntactic view, the semantic view, etc.)? There is still no consensus on questions such as these about the basic commitments of perspectivism and what problems it faces.

Here is one example that illustrates the diversity of positions that have gone under the name “perspectivism”. According to Massimi (2012), perspectivism should be taken to be an epistemic view, “a view on *how* we form scientific knowledge, as

opposed to a view about *what sort of objects* we have scientific knowledge of.” In a later introductory article on perspectivism, she similarly states that perspectivism is “first and foremost an *epistemic view* about the nature of our scientific knowledge”, while she denies that it involves metaphysical commitments, saying that it “is not intended to be a metaphysical view about scientific facts being perspectival or natural kinds being relative to scientific perspectives” (Massimi 2017). Contrast this with Giere, who takes perspectivism to involve the claim that kinds are relative to perspectives, and that there are no natural kinds in any more substantive sense (Giere 2006: 84-88). Along similar lines, an explicitly metaphysical formulation of perspectivism is provided by Chakravartty, who treats it as involving either of two claims:

(P1) We have knowledge of perspectival facts only, because nonperspectival facts are beyond our epistemic grasp.

(P2) We have knowledge of perspectival facts only, because there are no non-perspectival facts to be known.

A non-perspectival fact about a target system is thus a proposition that is true, independently of any particular perspective one may take with respect to it; it is true across perspectives. A perspectival fact is a proposition that is only true from, or within, or relative to, a given perspective (or limited set thereof). (Chakravartty 2010)

For Massimi, perspectivism is a view about how science builds perspectival knowledge of an objective, non-perspectival world. Thus, Massimi rejects perspectival facts. For Chakravartty, perspectivism is the view that perspectival facts are the only kind of facts we have access to. Even at this stage, then, we can see that there are important differences in the way that the view is understood. I am reminded of the multiple definitions of Kuhn’s “paradigms” (Masterman 1970); and perhaps, given the apparent similarities between “perspectives” and “paradigms”, it should not be surprising that perspectivism inherits similar ambiguities as Kuhn’s view. In addition, no doubt part of the problem here is that Giere’s own position is, as will become clear in the subsequent discussion, something of a moving target. It is

therefore unsurprising that later work on perspectivism, using Giere as a springboard, has taken off in many different directions. As a result of this, my goal in this thesis is not to evaluate perspectivism in general. Instead, I will focus specifically on Giere's perspectivism, and my aim is to develop a form of perspectivism that is coherent and defensible, and that satisfies the criteria that Giere specified as constitutive of a perspectivist position. Ultimately, the position I defend may be quite different in the details from what Giere would have endorsed, but I do think that it is firmly a *perspectivist* position, in the sense that he understood "perspectivism". With that said, I will now outline this approach to perspectivism.

1.2. Giere's perspectivism

As noted, Giere intends perspectivism to be a type of limited realism. In contemporary philosophy of realism, "realism" is generally treated as involving at least three components (cf. Chakravartty 2017a). First, a metaphysical claim that there is a mind-independent world; second, a semantic claim that scientific theories are to be interpreted literally; and third, an epistemic claim that we are justified in believing that at least some scientific theories correctly describe the world beyond the observable. That is, at least some scientific theories deliver knowledge of mind-independent entities, processes, or structure beyond the observable. Most contemporary realists also endorse a fourth claim: the reason why we are justified in taking our best scientific theories to be true is that the truth of these theories is the best explanation for their success. Truth is taken to play an important explanatory role.

It is worth starting by considering exactly what type of realism perspectivism is opposed to. Giere locates perspectivism in opposition to "objectivist realism". He identifies three claims that constitute objectivist realism (2006: 4), which can be summarized as follows:

(OR1) Science involves the discovery of truths about the external world, that form a permanent part of science and will not be overturned by future scientific research.

(OR2) Many of the truths discovered by science are laws of nature.

(OR3) Science is cumulative and converges towards one true theory over time.

This raises a problem. The point of perspectivism is to find a *via media* between objectivist realism and various forms of antirealism. But given this characterization of objectivist realism, it might be argued that at this point, perspectivism is rather trivial. Many philosophers already defend realist positions that do not involve commitment to all these theses. For example, almost all realists would weaken (OR1) to say that at best we achieve *approximate* truth, and that any given claim, no matter how well confirmed, might be overturned if a future theory achieves a closer approximation to the truth. That is, since realists only claim the approximate truth of our best theories, they need not suppose that any of those theories will be held permanently. There may always be room for improvement. The two other claims are also extremely contentious among realists. Cartwright (1983) denies (OR2), combining an antirealist view of laws with a realism about causal capacities; and Sankey (2008) denies (OR3), arguing instead that realism is compatible with denying an inevitabilist view of scientific progress. Nor are these unusual positions. Over the last several decades, realists have proposed a number of weaker forms of realism, in response to antirealist attacks. In this context, perspectivism looks redundant.

A little later, Giere characterizes perspectivism as follows:

For a perspectival realist, the strongest claims a scientist can legitimately make are of a qualified, conditional form: "According to this highly confirmed theory (or reliable instrument), the world seems to be roughly such-and-such." There is no way legitimately to take the further objectivist step and declare unconditionally: "This theory provides us with a complete and literally correct picture of the world itself." (2006: 5-6)

This paragraph suggests another claim involved in objectivist realism:

(OR4) We have, or could have, a theory that provides a complete account of the world.

This is a very strong claim, and one that it is unlikely any realist would endorse. Clearly, we do not currently have a theory that provides a complete account of the world, and there does not seem to be much hope for constructing such a theory.

This paragraph also indicates a stronger claim, that perspectivism involves a kind of relativism, according to which we may only make qualified claims. This may constitute a genuine challenge to traditional realism. We will return to this challenge in later chapters. At first glance, however, even this claim might be thought to be somewhat trivial. When I look at an object in the world, inevitably I see it only from a particular perspective, and this perspective will reveal some aspects of the object while obscuring others. In a similar way, when scientists construct models and theories, these are perspectives in sense that they focus selectively on particular features of target systems. No model can tell us absolutely everything about a particular target; such a model, even if we could construct it, would be useless to us. Since every model is selective, we must recognize our models as being constructed from a particular perspective. To properly understand any model, we must know what properties it is designed to be a model of. None of this seems to amount of a challenge to any account of realism that has been defended in the literature.

In what way does perspectivism depart from traditional realism? In describing perspectivism as a “via media”, we might assume that perspectivism joins the litany of more limited realisms that have been proposed in recent philosophy of science: structural realism (Worrall 1989), entity realism (Hacking 1983), deployment realism (Kitcher 1993), semirealism (Chakravartty 2007), and so on. However, I suggest that this is a rather misleading way to frame the debate, and I think that perspectivism at least attempts to provide a more novel approach to the realism problem. Indeed, perspectivism is compatible with all these other limited realisms. It is also compatible with certain forms of antirealism such as constructive empiricism.

All these other positions that imposed limits on realism – structural realism, entity realism, deployment realism, semirealism, constructive empiricism – introduce limits

to realism by restricting the set of propositions that are taken to be true. Consider a set of propositions that would be accepted by almost all scientists working in a field that is mature and exhibits a variety of predictive and technological successes; this set may include propositions such as “an atom of oxygen contains eight protons”, “the Earth is approximately 4.5 billion years old”, and so on. The naïve realist holds that all these propositions are true (or at least approximately true): we accept the truth of our best theories – that is the end of it. The structural realist holds that only those propositions describing structure are true. The constructive empiricist holds that we can know only those propositions describing observables are true. And so on.

The perspectivist, on the other hand, is silent on precisely which propositions are true, and focuses instead on the question of what it is for models to relate to the world such that claims generated on the basis of those models can be said to be true or false. The perspectivist tries to understand the model-world relation in a different way to the standard realist. If this is right, then one could be a perspectival structural realist, a perspectival entity realist; indeed, one could even be a perspectival empiricist. Take van Fraassen’s (1980) constructive empiricism, according to which science aims only for empirical adequacy. As we will see shortly, Giere suggests that we take our knowledge of colours as the exemplar of perspectival knowledge. Colours are, of course, observable properties. Perspectivism is fundamentally concerned with issues arising around truth, representation, facts, and so on. With this in mind, a condition of adequacy for a perspectivist theory is that it needs to be compatible with these different approaches to realism.

I have characterized this in terms of propositions, but this is perhaps misleading. Giere endorses the semantic view of theories and treats theories as sets of models rather than sets of propositions. But we can make the same point in terms of the semantic view. The perspectivist is proposing an alternative account of the way in which our theories and models represent the world. When we take a scientific proposition to be true, or when we take a scientific model to be accurate, what does this amount to? What is the proposition-to-world or model-to-world relation? In the next section I will discuss why perspectivism should be taken in this way, and I will outline in more detail just what the perspectivist alternative amounts to.

1.3. *The colour analogy*

For the perspectivist, Giere says, the strongest claims we can make are conditional: according to this model or theory, the world seems to be such-and-such. Models and theories are partly constitutive of *perspectives*. All our claims about the world are claims made from a particular perspective, and there is no way of stepping outside of all perspectives, no way of evaluating perspectives from a neutral, non-perspectival position. What exactly is a perspective? A useful general definition is given by Massimi, who states that a scientific perspective is:

the actual – historically and intellectually situated – scientific practice of any real scientific community at any given historical time. I understand scientific practice broadly to include: (i) the body of scientific knowledge claims advanced by the scientific community at the time; (ii) the experimental, theoretical, and technological resources available to the scientific community at the time to reliably make those scientific knowledge claims; and (iii) second-order (methodological-epistemic) claims that can justify the scientific knowledge claims so advanced (Massimi 2016a: 2)

This accounts for the many perspectivalisms noted in the opening section. Given such a broad definition of “perspective”, there are many different aspects of perspectives, and so many different things we might take facts or knowledge to be dependent upon. Moreover, perspectives occur beyond the sciences; not all perspectives are scientific perspectives. For example, Giere describes perspectives in the context of cartography:

I would like to say that the cultural background, the conventions for mapmaking, the designation of the region mapped, the specification of what features are mapped, and the degree of accuracy all determine a *perspective* from which the region is mapped. Every map reflects a perspective on the region mapped, a perspective built in by the mapmakers. In short, mapmaking and map using are *perspectival*. (Giere 2006: 75)

The most significant non-scientific perspectives for Giere are colour perspectives. Giere develops scientific perspectivism through an analogy to colour vision, which he argues “provides the best exemplar I know for the kind of perspectivism that characterizes modern science” (2006: 17). I will elaborate on the details of Giere’s position in later chapters, but it is worth stating the colour analogy from the start, as this provides a useful outline of the kind of view that Giere is aiming to defend. He begins with a description of colour vision. The human visual system responds to electromagnetic radiation in the range of about 390-700nm. It is sometimes assumed that colours are reducible to wavelengths of light, so that red, say, just is light in the range of about 640-700nm – call this view “wavelength reductionism”. This is a misleading oversimplification. There are two features of contemporary colour science that refute this simple reductionist view and favour a perspectival approach to colour.

First, there is the trichromatic theory, developed in the 1800s by Thomas Young and Hermann von Helmholtz, which claims that all the colours perceived by humans can be produced by different combinations of only three monochromatic lights: red, blue, and green. Mixing red and blue lights creates purple light; mixing red and green lights creates yellow light. Notably, the composite yellow formed from red and green lights (say, of 690nm and 530nm respectively) is perceptually indistinguishable from a spectral yellow produced by a single wavelength (say, 570nm), provided the mixture is just right. This is an example of metamerism, where two different illuminants appear to be the same colour. Similarly, there are reflective metamers, where two objects reflect different wavelengths of light but appear the same colour.

This occurs because the human retina contains three different types of cone cells, S-, M-, and L-cones, each sensitive to different wavelengths of light. S-cones are maximally responsive to about 430nm light; M-cones maximally responsive to about 530nm light; and L-cones maximally responsive to about 560nm light. For this reason they are sometimes thought of as the “blue”, “green”, and “red” cones respectively. However, each cone alone is colourblind. Simply stimulating the L-cone, for example, would not produce any colour experience. This is because a photon of light absorbed by a cone has the same effect on the cone no matter what the wavelength of the light is. A cone will produce a stronger signal either when the intensity of the light increases, or when the wavelength is closer to the cone’s peak

sensitivity; either way, the number of photons absorbed increases. In other words, with very weak light near the cone's peak sensitivity, there are relatively few photons arriving at the cone, but a high proportion of them are absorbed; while with very intense light of a wavelength to which the cone is not so sensitive, many photons arrive at the cone but few of them are absorbed. The cone's response is the same in each case. Colour perception arises from processing the *differences in the activation* of the three cone types.

Take just the L-cone. Light of about 650nm (seen by us as red) will produce the same response as light of about 490nm (seen by us as blue). They both produce the same activation in the cone. Furthermore, intense light of 650nm, to which the cone is not so sensitive, will produce the same response as dim light of 575nm, near the cone's peak sensitivity. So, a single cone cannot distinguish wavelength from intensity. But now add the M-cone, which has a different peak sensitivity. Any given light will excite the M- and L- cones differently. The 490nm light will excite the M-cones much more than the L-cones; the 575nm light produces a stronger response from the L-cones; and the 650nm light barely activates the M-cones at all. These differences in activation are preserved no matter what the intensity of the light is. An organism with both M- and L-cones is therefore able to distinguish these wavelengths, and this is a kind of colour perception. With two cones, this organism would have dichromatic vision. Many mammals are dichromats, and various forms of colourblindness in humans arise from lacking one or more of the three cone cells. Conversely, many birds are tetrachromats, with four different types of cones. Wavelength reductionism fails to account for metamerism and for the fact that the same wavelength will produce a completely different colour experience in a different organism: a red-green colourblind organism or a tetrachromatic bird may detect 650nm light, but they won't see red.

The second central feature of modern colour science that raises trouble for wavelength reductionism is the opponent process theory. Trichromatic theory alone doesn't explain several puzzling features of colour perception. In particular, some colour combinations are impossible. We can mix red and blue to generate purple, but we can't get any colours by mixing red and green or yellow and blue. Why is this? The reason is that information from the three types of cone cell is transmitted to two

opponent process systems, the red-green system and the yellow-blue system. The L- and M-cones feed the red-green channel, which calculates the difference in activation between the two cones. When the channel is excited, red is perceived; when the channel is inhibited, green is perceived. You can never perceive a colour that combines red and green because the channel cannot simultaneously be excited and inhibited. For the yellow-blue channel, input from the L- and M-cones is summed, and the channel calculates the difference between (L-cones + M-cones) and the S-cones. The result: four opponent colours from three types of cone. So colour space has a unique structure imposed by the operation of the visual system that is not mirrored in the structure of wavelengths of light.

If wavelength reductionism is implausible, what exactly is colour? Giere favours a perspectival account, and I will now outline six general features of colour perspectives, all of which, Giere argues, are also true of scientific instruments and scientific theories. I will make some brief comments here on how perspectivism is supposed to apply to theories here, but consider this point in much more detail later. I should note that my aim here is not to list six claims that define perspectivism. Rather, this is intended to provide a general impression of Giere's views and motivations. A fuller characterisation of perspectivism will be provided in later chapters.

1.4. Six features of colour perspectives

First, colour perspectivism is a form of colour interactionism. Colours do not exist independently of visual perspectives. A world without organisms would be a world without colours. Greenness does not somehow inhere in the grass independently of any observers. As Giere puts it, "Colours are the product of an interaction between aspects of the environment and the evolved human visual system" (2006: 31-32). Colours are not objective properties and colour perception does not simply present the world as it really is. There are two reasons for this. First, other organisms perceive colours in ways completely differently to us; second, the same colour can have a variety of physical causes while different colours can have very similar physical causes. So, even those who defend realism about colour, such as Byrne and Hilbert (2003), hold that colours are "anthropocentric" properties: colours are

reducible to objective physical properties, but only with reference humans. The colour red, for example, is the disjunctive set of surface spectral reflectances $\{X_1, X_2, \dots, X_n\}$, where these are grouped together on the basis of how they interact with the human visual system. For a group of extraterrestrial scientists who have never developed colour vision, this set would seem totally heterogeneous and of no relevance whatsoever to the sciences. What do all the elements of this set have in common? Just that they produce a particular kind of sensation in human perceivers.

A statement such as “grass is green” can be made only given the perspective provided by the human visual system. We cannot describe colours except with reference to particular visual systems. The visual system is not a mere window onto the world, but performs a great deal of sophisticated processing of the information it receives. Thus, we don’t simply see the world as it really is; we see the world from a coloured perspective. Similarly, Giere argues that scientific claims are always claims made from a particular theoretical perspective. While different theoretical perspectives are available, it is impossible ever to adopt a non-perspectival “view from nowhere”. Recall the quote from Giere given in Section 1.2 that for perspectival realists, the strongest claims we can make are of the conditional form: “According to this highly confirmed theory (or reliable instrument), the world seems to be such-and-such.” (Giere 2006: 5-6).

Second, claims about colour made within a particular perspective can be true or false. For a normal human trichromat, it is true that grass is green and that ripe lemons are yellow. To say that grass is blue, that grass is the same colour as the sky, is under normal circumstances just false, because this simply is not the view of the world that our visual perspective generates. Similarly, scientific claims are true or false relative to a theoretical perspective. Or more precisely, Giere’s view is that we can use a theoretical perspective to construct models that can be more or less similar to systems in the world. Exactly how this works, according to Giere, will be explained in more detail in Chapter 4.

Third, colour perception tracks genuine regularities in the environment. As noted above, most mammals are dichromatic; they are red-green colourblind. Why did red-green discrimination evolve in our ancestors? One popular theory is as follows: Our

ancestors were primarily tree dwellers. In dichromatic vision, everything in a forest looks mostly the same. But if you add red-green discrimination, ripe fruits pop out very prominently against the green leaves. Obviously it is of great advantage to any organism to be able to more easily detect food. Whether or not this evolutionary story is correct, it drives home the point that even if colours are in some sense subjective, colour discrimination does track objective facts. Colours should not be interpreted as simply a massive illusion; perceiving the colours of material objects is not like, for instance, experiencing an afterimage. Colours arise from the interaction of the world and our visual systems. Different visual systems provide different perspectives: these different perspectives are not inconsistent but simply track genuine regularities in different ways.

If, as Giere suggests, colour perception tracks genuine regularities, it seems that colour perspectivism still allows us to justifiably believe in certain non-perspectival properties. “The leaves are green” is dependent on features of human visual systems, but the distinction between the green leaves and the red fruit, the fact that we are tracking some relevant difference in the world, this is not obviously perspectival. A natural position might be that colours are perspectival properties, but we use colour to track non-perspectival properties. Giere’s own position on this point is unclear; as he says, “the typical human experiences the world from a coloured perspective” (2006: 32): it is compatible with this that what we are experiencing from our unique perspective may be non-perspectival properties. Whatever Giere would say about this, it’s worth noting that perspectivism doesn’t necessarily commit you to postulating perspectival properties. A perspectivist might hold that there is a mind-independent, objective world, and our visual systems, instruments, theories, etc., provide perspectives on this world. This was, of course, a key difference between Massimi and Chakravarty’s interpretations of the perspectivist position. I will return to this point later.

More generally, in viewing perspectivism as a *via media* between realism and social constructivism, I take it that for Giere one goal of perspectivism is to accommodate the social constructivist idea that truth is in some sense dependent on one’s perspective, while also including the realist claim that these are perspectives on the same world. Giere explicitly endorses the adoption of the methodological maxim that

the world has a unique causal structure (2006: 34-35). Having said this, we shall see later that embracing non-perspectival facts poses a serious threat to the philosophical significance of perspectivism. For now, I will simply note that this is a point on which perspectivists may differ.

Fourth, perspectives are partial. The visual system is responsive only to a certain range of electromagnetic radiation. Perspectives have limited inputs. Obviously the same is true for scientific instruments. Theories and models can also be seen as partial, in two respects. First, they only ever attempt to describe particular systems in the world, or they attempt to describe only processes at particular scales; a literal theory of everything would be far beyond our capacities – recall the Borges story of the people who wish to create a perfect map of their empire, with the result that they simply build an exact replica of the empire (Borges 1946). Second, theories and models focus only on particular properties of the target systems. A model of planetary orbits may simply ignore the atmospheric compositions of the planets.

Fifth, partly due to the partiality of perspectives, different visual perspectives – dichromatic, trichromatic, tetrachromatic – are just different, and should not necessarily be treated as incorrect or inaccurate. Where the trichromatic human sees a rug with a pattern of red and green, the dichromatic dog sees a uniform colour. These different colour perspectives are compatible (2006: 33): it is not that the pattern is “really” red and green, so that the dichromat misrepresents the real colour of the rug; rather, it’s that the trichromat can make finer discriminations in this respect. If the trichromat were able to communicate what she sees to the dichromat, the dichromat would learn something about the structure of the world that is consistent with what she already knows. Similarly then, different perspectives can be generated in the sciences and we can adopt a variety of different perspectives towards the same system. Perspectivism is compatible with a theoretical pluralism.

It’s worth noting that the partiality of perspectives perhaps suggests a way for the perspectivist to avoid more extreme forms of relativism. As Wimsatt (2007: 164) writes, “Traditional relativists have tried to recognize truths in different perspectives while in effect assuming they are each complete and all competing. They then seem inexorably pushed into saying that if these can each be right, then (in a sense)

nothing's right, and (therefore) anything goes." But perspectives for the perspectivist are not complete and all competing; various perspectives are required to understand the world.

Sixth, although different perspectives are not necessarily incorrect, different perspectives can be better or worse depending on one's purposes. This is because all perspectives are partial and imperfect. The visual system responds to a minor part of the electromagnetic spectrum, and often cannot discriminate different types of wavelengths (this is what leads to metamerism). Note that simply adding more photoreceptors might provide more information about colour differences, but it would have severe costs, primarily because adding new types of photoreceptors results in a lower maximum resolution. The majority of organisms with a visual system have a very limited number of photoreceptors, only around two to four in most sophisticated visual systems. For most organisms, there is simply no need to distinguish so many different types of colours: higher resolution is much more important.

One way to put this point is that perspectives exhibit *interestedness* (Chirimuuta 2016). Recall the example given above concerning how the human visual system arose in response to specific needs. Perhaps our forest-dwelling ancestors developed red-green discrimination to more easily distinguish the ripe fruit from green leaves; similarly, one possible reason why bees developed photoreceptors sensitive to UV light is that this makes flowers stand out much more prominently. While the trichromat can make finer discriminations than the dichromat, the dichromat may have a greater number of rods, permitting superior night vision, and this is possibly why so many mammals are dichromatic. We cannot fully understand a colour perspective, and the way it relates to the world, until we understand the context in which it is situated. Similarly then, "the theories and models of science are shaped by the needs and goals of the scientific community and wider society" (Chirimuuta 2016: 748). Scientific perspectives are evaluated not simply on the basis of whether they present the facts about the world, but also based on whether they satisfy the specific goals of scientific communities.

I hope that the preceding comments provide a sense of what colour perspectivism claims, and also how perspectivism might be extended to scientific knowledge. Giere

is clear that the discussion of colour is intended as an analogy for illustrative purposes, and not as an argument for scientific perspectivism. As we saw in Section 1.3, Giere describes colour vision only as the “best exemplar” for scientific perspectivism, and he explicitly says that a perspectival account of colour does not in itself support a perspectival account of instruments or theories, which require separate arguments (2006: 40). Nevertheless, given how much space Giere devotes to colour perspectivism, it is worth making a couple of comments about the analogy.

The most obvious point is that the analogy fails as a clear, uncontroversial example of perspectivism since there is much debate about whether the perspectivist account of colour is correct – and more generally, whether interactionist accounts of colour are correct. Many philosophers and colour scientists favour an antirealist or subjectivist approach to colour (e.g. Hardin 1988). Granted, Giere is right that objectivist realism does not seem to be plausible given the findings of modern colour science, but it’s notable that even most colour scientists accept this (though see Churchland 2007 for an attempt to defend colour objectivism from a naturalist framework). The debate among scientists and naturalistic philosophers is mostly between colour interactionists and colour subjectivists. This contrasts sharply with the attitude that many scientists have towards electrons, mitochondria, the citric acid cycle, black holes, etc. Is this a problem for the perspectivist? It’s natural to wonder, if perspectivism is right as a general account of scientific knowledge, why are so few scientists of most disciplines attracted to perspectivism, while colour scientists are sympathetic to colour perspectivism? In response to this, bear in mind that even those who are antirealist about colour will often use colour language uncritically. When a person says, “look at that red car”, she is simply reporting what she sees, and her statement is acceptable just as long as the car would indeed be classified as red by a majority of people. But when we “step back”, as it were, and consider colour in more detail just as colour scientists, we are pushed away from a realist account of colour. Similarly, once a theory is well-established, scientists can use its terms in an uncritical way. Only when stepping back and looking at the sciences “from the outside” does the perspectival nature of theorizing become apparent.

1.5. *Escape from Perspective*

I now turn to what I take to be the central challenge to perspectivism, a challenge that I will call “Escape from Perspective”. This challenge is raised by Giere, after he finishes his discussion of colour:

From what perspective do I make all the above claims about the science of colour vision? The simple answer is: from the perspective of colour science. ... The trouble with this simple answer is that it invites a further question. Even if it granted that colour vision is perspectival, that does not show that all the sciences involved in colour science are perspectival. In particular, it does not provide any reason to reject an objectivist understanding of these sciences. (2006: 40)

The problem here should be clear. Suppose that colour science tracks non-perspectival facts. In that case, it seems that we have a straightforward realism. Perhaps propositions about colour invoke an implicit reference to a particular visual perspective. “The sky is blue” is true, relative to the normal human visual perspective. But then we can use colour science, which describes the various features of the human visual perspective and the ways in which light interacts with objects in the world, to provide a non-perspectival, objective picture of the properties that are tracked by the human visual perspective. Apparently perspectival facts about colour properties are reduced to non-perspectival facts concerning surface reflectances, light, human visual processing, and so on. We have stepped outside the human visual perspective, to describe both the perspective and the objects and properties that it provides a perspective on. Hence why I call this the challenge of Escape from Perspective. This challenge has been raised by a number of perspectivism’s critics. Here is Votsis:

It is of course true that from the perspective of a normal human trichromat the sky appears blue. But *this* truth does not change when considered from the perspective of everyone else, i.e. normal non-humans and normal human non-trichromats, who is cognitively competent with respect to the given task. (Votsis 2012: 92).

Colour perspectivism is not particularly interesting if we can fall back on an objectivist understanding of colour science, and provide a non-perspectival account of the constituents of colour perspectives. Similarly, Chakravartty writes:

From the perspective I had of Peter over lunch in the Senior Common Room, he seemed a fairly tall man, but as I saw him in the distance some time after parting, he seemed rather small. This sort of perspectivism is uncontroversial because there are non-perspectival facts of the matter about the dimensions of Peter in our inertial reference frame that, in conjunction with facts about optics and my visual sensory apparatus, underwrite the differences in the appearance of his size. There is a height that he is, and then many ways he may appear to be from different perspectives. (Chakravartty 2010: 406)

In general, consider some supposedly perspectival property or perspectival fact, call it F . To say that this property or fact is perspectival is to say that it obtains from the point of view of a given perspective, call it P . In itself however, this is of little philosophical interest. We can think of the perspectival fact F as consisting in the way that some system F^* appears from perspective P . Then the claim will be that we can, in principle, describe F^* and P , and how it is that the interaction of F^* and P produce the appearance that F . There is a non-perspectival fact of the matter about F^* , P , and how their interaction produces particular appearances. Of course, we may not yet be able to fill in all the details. But this hardly supports any general perspectivist view.

What makes this challenge particularly pressing, in the context of philosophy of science, is that it is often thought of as a key feature of scientific inquiry that it achieves Escape from Perspective in this sense. The scientific description of the world is, as Nagel (1986) famously put it, a view from nowhere: a view from no perspective in particular. Human perception is perspectival; but our perceptual data is a product of mind-independent objects and properties that are separate from any of the perspectival appearances that they cause. More generally, any plausible view of science needs to account for its tendency to favour descriptions that are independent of particular human concerns. The challenge then is that any purported

perspectival fact can be reduced, in principle, to the non-perspectival, and it seems that the sciences are the best tool for achieving this reduction.

Escape from Perspective will lie in the background of much of the rest of this thesis, and one of the primary goals will be to develop a form of perspectivism that is insulated from this challenge. With this point in mind, it should be clear that perspectivism will be most plausible as a global thesis. If the perspectivist grants that there are non-perspectival facts in particular domains, she is at risk of simply capitulating entirely to the traditional realist. After all, nobody is a realist about everything; indeed, as we have already noted, most contemporary realists endorse one or another form of limited realism. If there are non-perspectival facts about, say, surface reflectances, it seems that the philosophical significance colour perspectivism is seriously undermined.

1.6. Historical parallels

In the history of philosophy, there has long been a tension between arguments appealing to the purportedly perspectival nature of of a particular domain, and realist defences that fall back on non-perspectival facts that underlie that domain. Consider the distinction between primary and secondary qualities, and the arguments against the objectivity of secondary qualities. For example, colours change depending on lighting conditions: the apple appears green under sunlight; but black under red light; hence the colour is not in the object itself. Or take sensations of heat and cold. Place the left hand in hot water, the right in cold water: now place them both in lukewarm water. To the left hand it will feel cold; to the right hand warm. In these cases, the same object appears to have different properties from different points of view. We should infer that the apple only appears green; the water only appears warm; these properties are not intrinsic to the object but a product of our reaction to the object.

There have been different ways of thinking about secondary qualities. Perhaps they are simply subjective; as Galileo said: "Hence if the living creature were removed, all these qualities would be wiped away and annihilated." Secondary qualities are simply features of consciousness, of how we react to the world. Alternatively, there is Locke's view that both primary and secondary qualities are powers in the object to

cause particular sensation in normal observers under normal conditions, where representations of primary qualities resemble the qualities in the object, but representations of secondary qualities do not resemble the qualities in the object. Today we might say: the quale of squareness resembles the shape of the object; the quale of redness does not resemble anything in the surface of the object, or anything in wavelengths of light, etc. Whatever exactly our account of secondary qualities, they turn out not to be robustly mind-independent. But, together with an appropriate account of the function of the human perceptual system, we can ground them in the mind-independent primary qualities.

In the history of the debate about primary and secondary qualities there is an interesting shift, noted by Ross (2015). In early modern philosophy, primary qualities were treated as explanatorily fundamental. Questions of realism are less important in this context. The point is rather that we explain secondary qualities in terms of primary qualities. In modern terms, this is a kind of reductionism: secondary qualities are reducible to primary qualities. Ross suggests that the distinction was thought of in terms of mechanical and non-mechanical qualities, with the former being the primary qualities. The mechanical qualities are the fundamental qualities of matter and all other qualities are to be explained in terms of mechanical qualities. Hence magnetism and electricity were classed as secondary qualities, as it was considered necessary to explain these phenomena in terms of more fundamental properties of size, shape, motion of particles, and so on. This distinction concerns the foundations of sciences, of what is fundamental to a scientific account of the world.

In more modern philosophy, the distinction between primary and secondary qualities became about *perceiver dependence*: in this approach, primary qualities are metaphysically perceiver independent; secondary qualities are metaphysically perceiver dependent. This is more closely connected to the concerns that tend to motivate perspectivism. And this distinction continues to be important. In philosophy of colour, colour is often treated as a relational property, as Giere treats it; but this contrasts with properties such as shape, which are thought of as independent of observers. However, notice that this changes the shape of the debate. Nobody would argue that colour is fundamental, in the sense that other properties must be explained in terms of colour properties. But we can hold that colour is a perceiver-

independent property, if a colour is taken to be simply a set of surface reflectances. Of course, even colour realists such as Byrne and Hilbert are forced to grant that colours are “anthropomorphic” properties, in the sense that we must define them in terms of the reactions of human perceivers; but on their view, colour track regularities in the world that are perceiver-independent.

The discussion of secondary properties raises a further concern about the colour analogy. Giere often talks as if “colour perspectivism” amounts to the view that colour properties are perspectival. A natural interpretation of this is that perspectival properties are secondary properties. Extending the analogy to scientific theories more generally, it is tempting to suppose that the scientific perspectivist is claiming that the properties examined in the sciences are also secondary qualities in a similar sense. There are two concerns about this.

The first is that the introduction of secondary qualities implies primary qualities that serve as their metaphysical ground. This immediately raises the problem of Escape from Perspective. Once we describe the primary qualities, we have a non-perspectival description. So this way of understanding perspectivism supports only a limited perspectivism.

Second, notice that this “secondary-quality perspectivism” would be a metaphysical thesis, a thesis about the ontology of the world. But recall the motivation for perspectivism. Giere is interested in developing a “middle path” between realism and antirealism. However, secondary-quality perspectivism is not an alternative to standard realism; it is simply realism with a rather unusual metaphysics. After all, the arguments for perspectivism are drawn from the sciences; Giere is explicitly naturalistic about philosophical methodology. So if we interpret “perspectivism” as secondary-quality perspectivism, then Giere has to be interpreted as taking at face value what he thinks our best science tells us about the nature and structure of the world: specifically, it tells us (according to Giere on this interpretation) that all properties are to be understood as secondary qualities. This is scientific realism of a *stronger* form than is usually defended in the literature. The secondary-quality perspectivist takes it that on the basis of contemporary physics, we are justified not just in believing that there are electrons, but also in believing a controversial

metaphysical theory about electrons. So secondary-qualities perspectivism is both implausible and at odds with Giere's stated goals. I take it that part of the reason why Giere asserts "the uniqueness of the world as a methodological maxim" (2006: 34-35), which involves acting on the assumption that the world has a single, unique causal structure, is precisely to distance his perspectivism from this interpretation. Giere is clear that uniqueness is only a methodologically useful, and he is explicitly agnostic about the metaphysics.

With these points in mind, the perspectivist must be cautious about how she uses the colour analogy. Either the perspectivist must repudiate the notion of "perspectival properties", and develop perspectivism in a way that does not make controversial metaphysical commitments (this is Massimi's approach); or she must understand perspectival properties or perspectival facts in a different way to the usual notion of secondary qualities. Ultimately, the approach developed in this thesis is of the latter sort, though the focus will be on the notion of perspectival truth rather than perspectival properties.

Similar considerations as lead to the (perceiver dependence) primary/secondary quality distinction also motivate the retreat to structural realist views. Traditionally, a common argument for structuralism is that we don't know the nature of things, but only their structural relations; and these structural relations do not change with perspective, or at least are not dependent on perspective in the same kind of way. This is the line taken in Russell's (1912) structuralism (for a useful summary see Frigg and Votsis 2011). Russell advocates a causal theory of perception: sense-data are caused by processes in the external world, and in this way they provide information about the world. We cannot assume that the intrinsic natures of anything in the world matches the intrinsic natures of sense-data, partly because of how appearances of things change in different conditions. That is, from different perspectives, objects will generate different sense-data. Frigg and Votsis suggest that Russell uses two principles to draw inferences about the world: first, the Helmholtz-Weyl Principle: "Different effects (i.e. percepts) imply different causes (i.e. stimuli/physical objects)"; second, the Mirroring Relations Principles: "Relations between percepts correspond to relations between their non-perceptual causes in a manner which preserves their logico-mathematical properties." (2011: 235)

Thus, Russell appealed to perspectival features of experience both to argue for sense data and against direct realism as a theory of perception, and to argue that we have only structural knowledge of the world. For Russell, both unobservables and observables are on a par: we don't perceive chairs, and we are only justified in believing the structural content of claims about chairs. Russell held that we have justified beliefs about relations of physical objects because these relations can correspond to the relations between sense-data. In the case of sense-data, we know the relational features and the intrinsic natures. By relational features Russell has in mind logical properties, temporal relations, spatial relations – or perhaps just mathematical properties; Russell (1927: 270) argues for total agnosticism about the external world beyond its mathematical properties. Our knowledge of relations must also be substantially limited: we don't know a relation like "darker than", because that would require knowing the intrinsic natures of the surfaces of objects. So in Russell's structuralism, we see both the appeal to perspectivity to undermine realism, and then the fall back on certain non-perspectival facts – in Russell's case, structural relations – to ground a limited access to the mind-independent world.

These manoeuvres – the primary/secondary quality distinction, and the retreat to purely structural realism – are attempts to specify a kind of realism that is insulated from the problems of perspective. Both of these are attempts to Escape from Perspective. In both cases, the attempt is of questionable success: arguably, perspectivity generalizes to the purported non-perspectival facts. The primary/secondary quality distinction faced the problem that primary qualities appear to share the same perspectival features as secondary quality. Just as colours change under different viewing conditions, so does shape, when a coin appears both circular and elliptical depending on the angle from which it is viewed. With respect to structural relations, Frigg (2006: 57-58) notes a similar problem:

a target system does not have a unique structure; depending on how we describe the system it exhibits different, non-isomorphic structures. If a system is to have a structure it has to be made up of individuals and relations. But the physical world does not come sliced up with the pieces having labels on their sleeves saying 'this is an individual' or 'this is a

relation'... Because different conceptualisations may result in different structures there is no such thing as the one and only structure of a system

Yet in the face of perspectival variation, there always seems to be the option of stepping back into a broader theory. We might say with respect to shape that we can specify certain viewing conditions, specify a certain perspective, from which we can determine by the laws of optics that the circle will appear as an ellipse; similarly, we can in principle specify conditions under which the green apple will appear black. For both shape and colour, there is perceptual variation; and for both shape and colour, the application of background theory creates stability; it creates a unified view of the world into which the many perspectives can be slotted, and understood in an apparently non-perspectival way. Can perspectivism be developed in such a way as to block this move? Clearly, it will do no good to simply point to some further, purportedly perspectival property. This always leaves it open to the traditional realist to Escape from Perspective by appeal to a background theory.

1.7. Outline of the thesis

The rest of the thesis is structured as follows. Part II discusses instrumental perspectivism, Giere's attempt to extend the colour analogy to the case of scientific instruments. I distinguish two arguments for instrumental perspectivism: first, a direct argument that attempts to find particular properties of instruments that render them perspectival; and second, an indirect argument which conjoins a perspectival account of perception with the view that instruments literally extend perception. Ultimately, I will argue that both of these routes to instrumental perspectivism fail. Instrumental perspectivism cannot stand alone: it is dependent on theoretical perspectivism, a perspectivist account of theories and models.

This leads us to Part III, which examines theoretical perspectivism. I introduce theoretical perspectivism in Chapter 4, where I examine the standard argument in favour of it, the incompatible models argument, and the different forms of theoretical perspectivism that have been developed on the basis of this argument. Again, the results here are rather dim for the perspectivist. The incompatible models argument does not provide a convincing reason to abandon traditional realism; and in any

case, most forms of perspectivism currently defended in the literature turn out to collapse into realism. Only an explicitly relativist perspectivism is available as a distinctive position, but the incompatible argument does not provide a good justification for it.

I then turn to what I think is a more promising strategy for the perspectivist, the failure of fit argument. The argument focuses on classification, and turns on the point that any proposition about the world must assume a particular classification scheme for dividing up the world. Classification schemes are provided by perspectives, and there is no perspective-independent standpoint from which such schemes can be evaluated. According to the failure of fit argument, traditional realism rests on the assumption that the world is “lexicon-dependent”; that there is a “language of the world” that the language of our theories might mirror. This argument is outlined in Chapter 5. Next, in Chapter 6, I present several potential responses to the challenge, and I argue that all of them fail. By far the most popular realist response is the appeal to “natural kinds”: the lexicon of the world is a lexicon of kinds, and the division between natural kinds is mirrored by the classification schemes of our best theories. The critique of the natural kinds response occupies Chapter 7.

Having dealt with these realist views, Part IV engages in the constructive project of developing a perspectivist view that respects the insights of the failure of fit argument. This occupies Chapter 8, which presents a perspectivist account of kinds and classification, drawing on Giere’s notion of “theoretical kinds”. This completes the positive case for perspectivism. After presenting this argument, I turn in Chapter 9 to the problem of Escape from Perspective, where I argue that the form of perspectivism proposed in this thesis has the resources to meet this challenge. In Chapter 10 I provide a summary of my view, and relate it to the six features of perspectives that were discussed in this chapter.

II

INSTRUMENTAL PERSPECTIVISM

Chapter 2

Partiality and Opacity

2.1. Introduction

In this chapter and the next, I will examine instrumental perspectivism. While much has been written on whether perspectivism provides a plausible analysis of models, laws, explanations, and various other activities involved in theorizing, Giere's perspectival account of instruments has so far received relatively little attention. I suspect that part of the reason for the lack of interest in instrumental perspectivism is that Giere himself presents it as a less controversial thesis as part of his book that ultimately aims to defend theoretical perspectivism. As we shall see later, it is not clear that instrumental perspectivism can be used in this way. In any case, instruments are interesting in their own right, and it is worth exploring the connections between instrumental perspectivism and theoretical perspectivism. I hope that this chapter makes some progress towards filling this gap in the perspectivist literature.

First, it is worth noting the overall structure of Giere's argument as it appears in his book. As Giere initially presents it, the case for instrumental perspectivism develops in three steps (for a summary, see Giere 2006: 14). First, he argues for a perspectival account of colour vision. Second, he takes it that perception in general is perspectival: "I will assume that the considerations suggesting that colour vision is perspectival can be extended to human perception more generally" (2006: 14), though no argument is given for the extension of perspectivism from colour vision to perception in general. Finally, he argues that the features that make colour vision perspectival are shared with scientific instruments: "observation using instruments is perspectival in roughly the same ways that normal human colour vision is perspectival" (2006: 41).

As I discussed in the previous chapter, the first claim is extremely controversial among philosophers of colour. There are a variety of other accounts of colour vision, and no consensus that a perspectival account is right. The second claim is also controversial, and surely one that Giere is not entitled to make without argument. It's not at all obvious that colour perception is relevantly similar to other forms of perception; much of the debate over primary and secondary qualities revolves around just this question. However, I am not sure that any of this is really a problem for Giere. In presenting an argument that builds from perceptual perspectivism to instrumental perspectivism, Giere is committing himself to much more than he really needs to, because the third step of his argument – detailing specific properties of instruments that render them perspectival – would, if correct, be sufficient in itself to establish instrumental perspectivism, regardless of how instruments relate to perception. Presumably the point of appealing to perception is to build from the less controversial claim (perceptual perspectivism) to the more controversial claim (instrumental perspectivism); but given the debates about the nature of perception, perceptual perspectivism is no less controversial than instrumental perspectivism.

In fact, Giere's discussion hints at two separate approaches instrumental perspectivism. The first way of thinking about instrumental perspectivism – call it the *direct argument* for instrumental perspectivism – simply points to specific properties of instruments that render them perspectival. For this argument, there is no need to ask whether instruments share any properties with perception. We just examine the properties of instruments directly. I think that this is the argument that Giere intends to endorse, hence why I say that his detour through perception ends up saddling him with controversial claims that he need not make. His discussion of perception should be taken more as a loose analogy to instruments, rather than as a crucial part of his argument. In fact, in his discussion of instruments, Giere specifies two properties that make instruments perspectival, *partiality* and *opacity*. In evaluating whether he is right, the nature of perception and its relation to instrumentation can be put to one side.

The second approach, which I will call the *indirect argument* for instrumental perspectivism, rests on two claims: first, that perception is perspectival, and second, that instruments literally extend perception. Instruments can, at least in certain

contexts, become part of the perceptual process. Therefore, instrumentation inherits the perspectivity of perception. Here the instrumental perspectivist appeals to the literature on the extended mind hypothesis. This argument is not endorsed by Giere, who explicitly rejects the extended mind hypothesis (2006: 110-112). But this argument suggests one way to apply Giere's perspectival account of perception in an argument for instrumental perspectivism, and therefore may be worth examining. In any case, I think that there are two separate routes to instrumental perspectivism. This is obscured in Giere's discussion because he seems to conflate two questions: What are the properties of instruments that make them perspectival? What is the relation between instruments and perception? In my discussion of instrumental perspectivism, I will try to keep these questions separate, and will focus on the direct argument and the indirect argument separately. The direct argument will be the topic of this chapter. The indirect argument will be the topic of the next.

The rest of this chapter is structured as follows. In Section 2.2, I discuss different types of instruments and present a general account of the kind of instrumentation relevant to instrumental perspectivism, drawing on Cartwright's (1999: 50) notion of nomological machines. In section 3.3, I examine Giere's direct argument for instrumental perspectivism, focusing on the properties of partiality and opacity. I find that Giere's argument does not justify the perspectivist account of instruments: partiality is trivial, while opacity is not well-defined in that it conflates two concepts, transformation and distortion. I conclude that the instrumental perspectivist who wants to forward a direct argument will need to defend the idea that instruments distort, a topic that is explored in sections 2.4 and 2.5. These sections develop the case against the direct argument, as it turns out that the direct argument makes instrumental perspectivism dependent on theoretical perspectivism in a way that robs it of philosophical significance, and that it also introduces a variety of puzzles about how perspectives should be defined and demarcated.

2.2. Instruments as causal processes

There are many types of instruments used in the sciences, and it may not be reasonable to expect a philosophical analysis of instruments to apply to all of them. We can distinguish three types of instruments: instruments that produce phenomena,

or *production instruments*; instruments that detect phenomena, or *detection instruments*; and instruments as analogies, or *analogical instruments*. My taxonomy here is inspired by Rom Harré (2003; 2010), though with some differences.

2.2.1. *Varieties of instruments*

Production instruments. These instruments are used to create or alter things, where what is created or altered by the instrument is the goal of using the instrument. Consider the use of instruments for the purposes of cloning, or the use of a thermal cycler in polymerase chain reaction, or radiotherapy to kill cancer. We use the instrument as a tool to achieve some practical goal.

Detection instruments. Here we use the instrument to detect some otherwise unobservable phenomenon, or to help correct our judgements about observable phenomena. Following Harré (2003) we can distinguish “primary” and “secondary” quality detection instruments, drawn from the traditional primary/secondary quality distinction. In primary quality detection instruments, the output of the instrument is taken to have relevant resemblances to the input: perhaps it is similar to the input, or it has the same structure as the input. An atomic force microscope is a primary quality detection instrument since this produces an image that maps the topography of the sample, showing areas of higher and lower elevation. Brighter areas of the image might map to higher areas of the sample; darker areas of the image map to lower areas. The output resembles the cause in a particular way. In contrast, in secondary quality detection instruments, the output does not resemble the cause. The movement of mercury in a thermometer is caused by but does not resemble the average kinetic energy of the molecules of a substance; the change in colour of litmus is caused by but does not resemble acidity.

Analogical instruments. These are physical models of systems in the world. Consider the use of scale models to study geological processes. Oreskes (2007) discusses Alphonse Favre’s attempts in the 1800s to test whether the patterns of mountain folding were produced by the contraction of the Earth. Favre created a model consisting of bands of clay placed along stretched rubber; as the rubber was released, the clay deformed. The patterns observed in the deformed clay were

similar to the patterns of rock layers in mountains, thus lending plausibility to the contraction hypothesis. An analogical instrument exhibits or realizes a particular type of causal process. Despite the numerous differences between a small model made of rubber and clay and the crust of the Earth the hope is that the clay model exhibits a relevantly similar kind of causal process so we can extrapolate from it to the Earth.

The boundaries between these types are not fixed: exactly how we classify a particular instrument will depend on how scientists use and interpret that instrument. When animals are used in space testing programs, the animal may become both a detection instruments (detecting the effects of G-forces on animal flesh, say) and an analogical instruments (the main goal for animal testing in this context is usually to use the animal as a model for the human). How we classify an instrument may also depend on our broader epistemology of science.

What types of instruments are relevant to instrumental perspectivism? Production instruments can be put aside in this context, as the output of such instruments is not taken to represent any systems under investigation in the world. All the examples of instruments discussed by Giere, such as gamma ray telescopes and magnetic resonance imaging, are primary quality detection instruments. The output of these instruments is presented in the form of images that we think resemble in some ways the properties of the system being studied; or at least the output of these instruments is used to construct such images. I assume that instrumental perspectivism can also be extended to secondary quality detection instruments. As for analogical instruments, we might speak of these as offering as “perspective” on phenomena in the sense that we are asked to view the phenomena as being *like* some other process. The clay and rubber model of the Earth’s crust encourages us to view the Earth’s crust as being like an instrument made of clay and rubber. However, analogical instruments are physical models, and so a perspectival account of these is better construed as an application of theoretical perspectivism. So in this chapter we will focus on detection instruments, and henceforth the term “instruments” will be used to refer specifically to detection instruments.

2.2.2. *Nomological machines*

When thinking about scientific instruments it is useful to consider what I am calling the *instrumental model*. The instrumental model is a model of the production of the instrument's output; this is similar to what Rothbart (2007) calls the "design plans" of instruments. A key point here is that instruments are special types of causal processes, and this causal process is described by the instrumental model.

Instruments are nomological machines in Cartwright's sense; Cartwright defines a nomological machine as "a fixed (enough) arrangement of components, or factors, with stable (enough) capacities that in the right sort of stable (enough) environment will, with repeated operation, give rise to the kind of regular behaviour that we represent in our scientific laws" (Cartwright 1999: 50). The instrumental model tells us how the processes going on in the world plus the functioning of the instrument produce a particular output. With the instrumental model we can draw conclusions about the world from the instrument's output. The output of the instrument is modelled as the end of a causal chain; and this allows us to use the output to infer something about an earlier stage in the causal chain or another object in the causal chain.

Let's take Giere's example of gamma ray observations of the Milky Way (Giere 2006: 45-48). Different types of telescopes are sensitive to different wavelengths of electromagnetic radiation. The Imaging Compton Telescope, or COMPTEL, uses Compton scattering to detect gamma rays. Compton scattering occurs when a photon collides with a charged particle such as an electron, causing the photon to be scattered and the electron to recoil. The photon loses energy to the electron, and the angle of the electron's recoil is determined by the amount of energy it receives.

Very briefly, the instrument works as follows. It consists of two detectors, each surrounded by photomultiplier tubes. Each tube is sensitive to gamma rays in the range of 1 to 30 MeV. The incoming gamma ray is Compton scattered in the first detector, losing energy; the recoiling electron produces a scintillation measured by the photomultiplier tubes. The gamma ray is then Compton scattered again in the second detector, again producing a scintillation measured by photomultiplier tubes. From the intensities of the two scintillations we can infer the energy of the original

gamma ray. The organization of the photomultiplier tubes also allows the direction of the gamma ray to be determined. The entire instrument is surrounded by a scintillator shield that detects any stray charged particles whose interaction with the detector can be discounted.

Scientists use the data generated by COMPTEL to construct images of the Milky Way. We understand these images, and can use them to draw conclusions about the Milky Way, only given the instrumental model and various background theories and assumptions – that is, we use these images to draw conclusions about the Milky Way only because we understand the production of gamma rays in the Milky Way, and we understand the way that gamma rays interact with other particles, etc. This is what allows us to conclude that a bright spot on an image produced by COMPTEL shows us an especially intense source of gamma radiation. The instrumental model is itself drawn, of course, from various background theories and models: the instrumental model of COMPTEL includes Compton scattering and scintillation, processes that are described and explained by physical theory.

The instrumental model is incomplete, or imprecise. In explaining how COMPTEL functions, we do not say exactly what its input and output will be. Instead we make only conditional claims: if its input is X, its output will be Y. The instrument's output allows us to precisify the instrumental model: by specifying the output of the instrument, it thereby specifies or at least narrows down the possible input. When we embed a particular output into the instrumental model, it makes one part of the instrumental model precise, and so also makes precise other parts of it. Of course, it cannot achieve a complete precisification. Measurement of the direction of a gamma ray may be affected by various factors such as atmospheric conditions. Hence from the output, we can infer the location of a source of gamma radiation only within a particular margin of error. The instrumental model itself often tells us the margins of error: thus the influence of atmospheric conditions is included in instrumental models of telescopes.

At this point a potential objection arises. The instrumental model is important because it allows us to interpret the terminus of a causal chain comprising a nomological machine. But we can of course perform further manipulations on an

instrument's output, and so create further outputs from it. It might seem that in some of these cases we can dispense with the instrumental model, since humans need not interpret the output of an instrument. Consider the many cases where scientists use computers to analyse the output of instruments and they never even see the raw data that is fed into the computer: Humphreys (2013: 65) discusses how in computerized tomography, we use computers to manipulate sinograms to draw out information that would not otherwise be accessible to a human. One of the consequences of the emergence of big data (Lyon 2014) is that it's simply impossible for scientists to interact with raw data from instruments.

How does the view outlined here account for this? I would say that in cases where a computer manipulates the data before it is interpreted by scientists, this can be viewed as simply an extended instrument, understood in terms of a new, extended instrumental model. Note that computer manipulation of data is only useful when the manipulations are well-defined and well-ordered so that we can rely on it to produce a particular type of output. To take a simple example: Suppose that instrument *X* produces output *a* and instrument *Y* produces output *b*. Here we have two outputs. It may then be useful to combine *a* and *b* to generate new data *c*. For example, we might combine an image of the Milky Way produced by COMPTEL with an image produced by an ultraviolet telescope – say, if our best theory of star formation entailed that stellar nurseries emitted strong radiation in specific gamma ray and ultraviolet wavelengths. Obviously, the combination of *a* and *b* into *c* could be done by a computer; in principle scientists need not interact with *a* and *b* at all. But to have any understanding of *c*, scientists must still appeal to the instrumental models of *X* and *Y*. If we have set up a system where a computer regularly generates output of type *c* from output of type *a* and type *b*, so that scientists can ignore output of type *a* and type *b*, then *X* + *Y* + computer can be thought of as a new instrument, and the instrumental model is a model of the production of output of type *c*. Computer manipulation of data is of course often much more sophisticated than simply combining images, but I think the same point applies. So, I don't think that computation challenges what I have said about instruments.

To summarize my view of instruments: Instruments are nomological machines, well-ordered causal processes, and we can draw conclusions about the world from our

use of instruments by appealing to an instrumental model that describes the causal connections between systems in the world and the output of the instrument. I hope that what I have said so far is not too controversial. I don't think that Giere would disagree with any of this. What then is the case for instrumental perspectivism, and what does this view add to our understanding of instruments?

2.3. *The direct argument: Partiality and opacity*

In this section we will examine Giere's argument for instrumental perspectivism, which proceeds by attempting to show that particular properties of instruments supports a perspectival analysis. Giere's basic intuition is that instruments do not simply reveal how the world is; rather, they reveal how the world is *from a particular point of view*. So Giere intends to defend a kind of qualified realism – a perspectival realism. More precisely, according to Giere, instruments are perspectival in two ways:

First, like the human visual system, instruments are sensitive only to a particular kind of input. They are, so to speak, blind to everything else. Second, no instrument is perfectly transparent. That is, the output is a function of both the input and the internal constitution of the instrument. (Giere 2006: 14)

I will call the first property *partiality*: instruments are partial in that they respond to only a certain type of input. I will call the second property *opacity*, as Giere describes it as the instrument failing to be “perfectly transparent”: instruments are opaque in that the output of any instrument depends on the input plus the instrument's internal processing.

Giere discusses several different instruments, but a single example will suffice to illustrate the main philosophical points, so we will take COMPTEL, already discussed earlier. It's not difficult to see how COMPTEL exhibits both partiality and opacity. COMPTEL is partial because it is sensitive only to a particular range of electromagnetic radiation. It cannot measure visible light or neutrino flux or sound. It is opaque because its output depends on processes internal to the instrument, such

as the structure of the photomultiplier tubes. Images constructed from the data provided by other types of detector would provide very different views of the Milky Way. What COMPTEL produces is not simply an image of the Milky Way, but an image of the Milky Way from a particular instrumental perspective; as Giere puts it: "Observation does not simply reveal the intensity and distribution of gamma rays coming from the centre of the Milky Way, it reveals the intensity and distribution of gamma rays as indicated by COMPTEL or OSSE or..." (2006: 48; emphasis in original). Later Giere says that we "cannot detach the description of the image from the perspective from which it was produced" (2006: 56). In analysing the output of any detector, we must consider the internal structure and processing of the detector itself.

There is no question that instruments are sensitive to a particular type of input, and that the output is dependent on the internal processing of the instrument. In the rest of this section we will consider what we should make of this fact. Do these properties challenge a realist account of instruments, or require the realist to qualify her realism in any way?

Partiality can be dealt with swiftly. Clearly COMPTEL is partial, as it is sensitive only to electromagnetic radiation within a certain energy range and can only detect gamma rays that reach the Earth; regions of denser gas and dust throughout the galaxy will obscure many parts the sky. Furthermore, analysis of the data may eliminate some of the detection events. Partiality is not a controversial thesis, and it's hard to see how it could support any particular philosophical analysis of instrumentation. You don't need to be an instrumental perspectivist to see that instruments are responsive only to certain types of input; both realists and antirealists will also accept this. One way to see the problem for the instrumental perspectivist position here is to ask: what is the alternative? What would it mean to reject partiality? If an instrument were not partial, then it would be sensitive to everything. Even if such a thing is logically coherent, there would obviously be no way to use it. Instruments are useful precisely because they are used to track only specific properties that we are interested in. Since everybody accepts partiality, instrumental perspectivism is trivial if it rests on this claim.

It might be argued that I am being uncharitable to the instrumental perspectivist here. Perhaps partiality is simply one premise of an argument for instrumental perspectivism, and it is after all a good thing for premises to be uncontroversial. But I don't see any such further argument from partiality to perspectivism in Giere, nor do I know how this argument might go. Giere simply cites partiality and opacity as features that make instruments perspectival. Since partiality is already accepted by everybody, we can put this feature to one side. Perspectivism as partiality is not an interesting form of perspectivism.

Let's turn then to opacity. This is the claim that the output of an instrument is dependent on the input plus the internal processing of the instrument. Now initially it may seem that this, like partiality, is a trivial claim. The images produced by COMPTEL depend on its internal processing; obviously, without any internal processing, it wouldn't produce any images at all. A different gamma-ray telescope, sensitive to the same wavelengths, may well produce different images. Imagine that aliens who see only in ultraviolet light develop their own COMPTEL telescope; the images this instrument produces would look different, since it has been constructed for the ultraviolet-sensitive eyes of the aliens.

So it may be thought that opacity is again trivial, in which case the philosophical significance of Giere's instrumental perspectivism is lost. But the perspectivist draws a controversial moral from the fact that the output of an instrument is dependent on its internal processing. One of the central ideas of instrumental perspectivism is that perspectives contribute to the content of an instrument's output, and therefore an instrument only ever shows the world *as it appears from that instrument's perspective*, not the world as it really is. This is one way in which perspectivism challenges standard realism: we never have access to the world, without qualifications; only the world from a particular perspective.

We can get a clearer idea of what opacity involves if we consider what it would take for an instrument *not* to exhibit opacity. The metaphor of transparency vs opacity is telling. Look at a tree, and then look at a tree through a window. The window is literally transparent. If we were to interpret the window as an instrument, we might say that the "output" of the window (what will be detected on the other side of the

window) depends only on the input, since the “internal processing” of the window makes no difference to how the scene outside appears to the eye. Of course, this is not quite true. There will always be some detectable differences between looking at a tree and looking at a tree through a window. But the information reaching your eyes in the latter situation may increasingly approximate the information reaching your eyes in the former. The key point is that in most contexts, a clear window does not transform the information in any significant way. The window-instrument exhibits *partiality*, since it allows only certain types of input through – electromagnetic radiation in the visible spectrum – but it is not *opaque*; it does not alter the information. Given the input, we could predict the nature of the “output” without considering the window at all. This is not the case for a gamma-ray telescope or an fMRI scanner. In that sense the window-instrument shows the world “as it really is”, whereas the output of the gamma-ray telescope and the fMRI machine is “conditioned” by these instruments.

However, there are at least two things that are very strange about the paragraph above. First, it does not seem right to say that the window shows the world as it really is. Speaking more carefully, the situation is that what we see when we look through the window is relevantly similar to what we see when the window is not there. The transparency of the window lies in the fact that it makes no difference to what is seen. This points to the second strange thing: attempting to interpret a window as any kind of instrument. Instruments are used to make our discriminations more reliable, or to detect properties that would otherwise be undetectable. But this requires that information otherwise hidden to us be transformed into something that we can perceive in some way. Under normal circumstances a window doesn’t detect anything. Light simply passes through it, and as a result we simply see through it to whatever is on the other side. For this reason, a window is totally useless as an instrument. The moral: to detect properties in the world, you have to alter the information about those properties. But now it seems that detection *requires* opacity! After all, no realist would want to say that instruments are literally like windows. All parties to the debate accept that instruments involve transformation of information.

Consider again the example of COMPTTEL. We can’t detect gamma rays. So the input must be altered by the instrument in order for it to be of any use to us. By

transforming the input from, say, a distant galaxy the instrument makes certain properties of that galaxy salient to us, properties that would be invisible to us were it not for the instrument. Should we say that the instrument is not revealing the world – or, in even stronger realist terms, the world as it really is? The only way I can make sense of such a claim would be that since the instrument transforms the input, it thereby distorts the input. In this sense the instrument contributes to the output.

In my view there is a serious confusion here. I think that much of the appeal of instrumental perspectivism rests on conflating *what* is detected with the *means by which* things are detected. The point is perhaps best seen by returning to the comparison with visual perception. We might say, speaking loosely, that we perceive light – but of course this is not true: we perceive objects, and light is the means by which we perceive objects. Similarly, the internal workings of the visual system are part of the means by which we perceive objects. Processes involving light, the lens of the eye, the retina, the optic nerve, the lateral geniculate body, etc., contribute to the content of visual representations insofar as we wouldn't have those representations without those processes occurring. In general, there is not a detection of X without some process occurring that constitutes the detection, i.e. without there being a vehicle for the detection.

With this in mind, what exactly is it that is being distorted by the instrument? The gamma-ray telescope allows us to detect properties from a distant galaxy by transforming a certain type of information from that galaxy. It makes no sense to view this transformation as any kind of distortion because the transforming processes are the means by which the properties of the galaxy are detected! Obviously these properties are not themselves distorted. We can't change the properties of a distant galaxy. There are some situations where detecting the properties of a system changes the system, but this isn't one of them, and in any case such situations are irrelevant to the instrumental perspectivist's point. To sum up: detection requires transformation. Transformation is not distortion (not in itself at least). What then does it mean to say that an instrument is opaque? If it means that the instrument transforms, then this is true but, like partiality, trivial. For an instrument to detect otherwise unobservable properties, there must be some processes occurring in the instrument that transform the information. If it means that

the instrument distorts, this is controversial but not adequately supported by Giere's argument. Giere's concept of opacity seems to combine both notions, of transformation and distortion. What are we to make of the notion of distortion? Is there some interesting sense in which instruments distort?

2.4. Do instruments distort?

According to Giere, instrumental perspectivism consists in the fact that all instruments are both partial and opaque. If the argument of the previous section is correct, then what is philosophically controversial in instrumental perspectivism is the idea that instruments *distort* information. This section therefore takes up the question, in what sense could an instrument distort information? Drawing on the comments of the previous sections, it might be argued that there is something misleading in describing the output of any instrument as either distorted or veridical. Instruments are well-ordered causal processes that transform information. On the basis of the output of the instrument we draw conclusions about the world that may or may not be correct, depending on whether we are applying the right instrumental model. It is only when we make inferences from the output that the question of distortion arises.

However, I think that this conclusion is too quick, and that the instrumental perspectivist can make sense of the idea that instruments distort. The key is that an object counts as an instrument – a detection instrument – only if it is used by people in the right kind of way: in particular, people must apply an instrumental model to it. Without the instrumental model, we cannot specify what the “input” or “output” of the instrument even are; we don't really have any data at all on the basis of which we might make any claims about the world. In other words, it is not that we have an instrument, with a particular type of output, and we then apply a model to draw a conclusion about the world. Rather, a nomological machine only becomes an instrument when it is used and understood in a special kind of way, which involves the application of an instrumental model. Put simply: there is no instrument without an instrumental model.

It is of course true that instruments sometimes create surprising phenomena that is not covered by any theory or model: Roentgen's fortuitous discovery of X-rays is a classic example here. But note that such surprising discoveries are surprising only against a background understanding of previously accepted theory. Initially, Roentgen was simply faced with a strange phenomenon that he couldn't explain: a faint green glow on a barium platinocyanide screen, produced while operating a cathode ray tube. Surely this didn't count as a discovery of X-rays, or as a perspective on X-rays, or anything like that. It took seven weeks of work before Roentgen had any idea what he was looking at. Per the taxonomy of section 2.2.1, all Roentgen had before some theoretical understanding had been achieved was a production instrument: set this equipment up in the right way, switch it on, and you get a faint green glow. This is a literal "engine of creation" in van Fraassen's (2008) sense. It's an instrument of some sort, but we have an *X-ray detection instrument* only when we can use the screen to take measurements of X-rays. In general, something is a detection instrument only when it is used to take measurements of systems, and that requires an instrumental model. The issue is that without an instrumental model, we can't link up the phenomena produced by a piece of equipment to other processes in the world. How do we know that we have an image of the Milky Way – or even just an image of the Milky Way *from the perspective of COMPTEL*? Only because of some model linking processes in the Milky Way to the functioning of the COMPTEL instrument.

In the literature on measurement, it is standard to distinguish four different types of measurement scale: nominal, ordinal, interval, and ratio (Stevens 1946; Tal 2013: 1164). In a nominal scale, objects are sorted into sets without any specific order, as when we sort people into male and female, or we sort people into different nationalities. An ordinal scale sorts the measurands into a particular rank order: the Mohs scale of hardness ranks minerals into ten categories based on whether one can scratch another. This scale is qualitative and no meaningful arithmetical operations can be performed on it, since there is no way to combine degrees of hardness. An interval scale has equal intervals, and so allows arithmetical operations on the intervals of the scale but not on specific values of the scale. Consider the Celsius temperature scale. We cannot say that 100°C is twice as hot as 50°C, because the zero point is arbitrary. However, the difference between 100°C and

50°C is the same as the difference between 50°C and 0°C. Finally, a ratio scale also has equal intervals but the zero point is not arbitrary. 100kg is twice 50kg, and 0kg is no mass.

We cannot make sense of a measurement of a particular quantity unless it is placed within a measurement scale. If I say that topaz has a hardness of 8 on the Mohs scale, this information is useless unless you know what kind of scale the Mohs scale is, and you know where various other items are placed on the scale. It almost goes without saying that no instrument can in itself specify a measurement scale. The type of scale used must reflect the relations between the objects being measured, and this imposes limits on how we might analyse data from an instrument. As Hand (2004: 73) says, “when regarding the numbers as measurements of some underlying property, only those operations may be performed which correspond to some empirical operation between the objects measured.” The instrumental model is what specifies this correspondence.

Furthermore, even with the measurement scale given, the instrument cannot itself give the value of the measurand. To make a measurement, we must fix the appropriate degree of precision. Collins (2010) gives the example of measuring height. Suppose I want to convert my height from inches to centimetres. If my height is 69 inches, and if I know that there are 2.54 centimetres to an inch, then I can calculate that my height is 175.26 centimetres. But Collins points out that convention does not allow me to say: “my height is 175.26 centimetres.” This is not an acceptable measurement report; we do not measure heights to the one hundredth of a millimetre, because we know that height varies by several hundredths of a millimetre even between a person’s breaths. Instead, I might say that my height is 175cm, or 175.3cm, or 175.5cm, depending on context. How do we decide what counts as an appropriate degree of precision? Among other things, we must consider the properties of the measurand and the margin of error of the instrument. The margin of error is given by the instrumental model.

Instruments are used to take measurements. Taking a measurement requires specifying a value within a particular scale, and to a particular degree of precision. The output of an instrument cannot simply reveal the property of an object. To draw

a conclusion about the object, we must locate the instrument's output in a logical space provided by the instrumental model and background theory (cf. Van Fraassen 2008: 164).

Am I overstating the dependence of instruments on theory? Many writers have argued that instruments are in some important sense independent of theory: Hacking (1983) urges that some experiments have "a life of their own", as it were, and can be conducted without guidance from theory. In particular, arguably we can know that an instrument is reliable without appealing to any sophisticated theory. But the key point here is *sophisticated theory*. Instrumental models need not be theoretically sophisticated, and much literature has examined the ways in which models are constructed relatively independently of theory (Morgan and Morrison 1999). You don't need to know the precise details about the nature of light to understand the instrumental model of the optical microscope: this model has been in use since the earliest days of microscopes, surviving through corpuscular theories, wave theories, the photon theory. Even Hacking accepts that to make any judgements about unobservable entities, we need to know various "home truths" about them, certain basic generalizations about important causal properties (Hacking 1983: 265). Such home truths are enough for the instrumental models of instruments. I take it that those who assert the independence of instruments from theory mean by the word "theory" something more like a complex, formalized system, intended to apply to a wide range of phenomena, partially built on universal laws.

There can be no detection instrument without an instrumental model. It is not quite correct that we have an instrument, with a particular type of output, and then we draw inferences about the world by applying the instrumental model to that output. The instrument and the instrumental model are more closely intertwined than that. To identify some event as the output of a detection instrument is already to apply an instrumental model to it. So there *is* a sense in which instruments can be said to distort, because the instrumental model may be mistaken in various ways. Recall the point made by Hand that the permitted operations on numbers of a measurement scale must correspond to some empirical operation between the measurands. We may be mistaken about what these empirical operations are. On this point then, the instrumental perspectivist is vindicated. Furthermore, there is some *prima facie*

plausibility in the idea that even our best instrumental models are subject to distortion, because as Levins (1966) argues, all models face a trade-off between accuracy, precision, and generality, and so there can be no perfect and complete model of any system. All models will feature limitations and idealizations. Whether the realist can accommodate this fact is beyond the scope of this chapter (see for example Shaffer 2012 for a defence of realism on this point); but it does provide some initial support to a perspectival account of all instrumental models.

However, this is a Pyrrhic victory for the instrumental perspectivist, because it follows from this that instrumental perspectivism is dependent on theoretical perspectivism. We can defend instrumental perspectivism only by defending a perspectival account of the instrumental models. The distortion in the instrument is due to inaccuracies in the instrumental model; if the instrumental model is correct, then the instrument simply tracks properties in the world and it would make no sense to treat its output as distorted. So on a straightforward realist interpretation of instrumental models, there can be no case for a perspectival analysis of instruments. One way to see this point is to note that even if we treat instruments as perspectival, a correct model of the internal processing of an instrument will allow us to “step outside” the instrument’s perspective. Here we have a straightforward application of Escape from Perspective. This point is nicely illustrated by how we use our theoretical understanding of how visual perspectives operate to escape the limitations of our own specific visual perspective, as in Chakravartty’s example, quoted in Section 1.5, that Peter’s height differs from different perspectives:

This sort of perspectivism is uncontroversial because there are non-perspectival facts of the matter about the dimensions of Peter in our inertial reference frame that, in conjunction with facts about optics and my visual sensory apparatus, underwrite the differences in the appearance of his size. There is a height that he is, and then many ways he may appear to be from different perspectives. (Chakravartty 2010: 406)

Recall Giere’s example of COMPTEL. Part of the intuitive appeal of instrumental perspectivism lies in the face that when we use COMPTEL to create an image of an object, such as an image of the Milky Way, the object will appear very different to

how it appears to the naked eye. So there is surely some sense of the word “perspective” in which it is true to say that this is an image of the Milky Way from the perspective of COMPTEL. But this difference simply results from the fact that COMPTEL is designed to reveal different properties to what is detected by the naked eye. Then what allows us to claim that we have an image of the Milky Way from the perspective of COMPTEL is some model of the functioning of COMPTEL and its relation to processes going on in the Milky Way. If this model is correct, then COMPTEL does not distort. In the same way, a realist would insist that the human visual system does not distort the size of Peter when the amount of the visual field that Peter occupies changes with changing distances. Indeed, it is precisely this feature of the visual system that allows us to estimate Peter’s true size! Put simply then, distortion resides in the instrumental model, so instrumental perspectivism requires a perspectival account of the instrumental models.

One might object at this point that I have overlooked other ways in which instruments can be said to distort. Consider phenomena such as spherical aberration, where the light rays from a lens do not all converge on one focal point. In this case, even if we had a completely correct instrumental model, and so could draw the right inferences about the input to the instrument, intuitively there is still some sense in which the optics distort. Isn’t spherical aberration a source of distortion, even when we accurately model the spherical aberration? After all, it’s natural to speak of using combinations of lenses to “correct” spherical aberration. One way to put this objection is that not all kinds of transforming processes are equal: it’s not enough simply to have a model which correctly describes how an instrument transforms the input; we also prefer particular kinds of transformation processes over others.

My response to this is that while I agree that it is intuitive to describe spherical aberration as a form of “distortion”, I am not sure that we are using the term in the sense that is relevant to the instrumental perspectivist. The “distortion” involved in spherical aberration is a practical matter. What exactly is the problem with spherical aberration? It is that spherical aberration reduces resolution and clarity, so optical instruments that are subject to spherical aberration are more difficult to work with, and worse, information about the input will be lost. Suppose we are observing a specimen using a microscope with severe spherical aberration; suppose also that

our instrumental model of the microscope models the aberration. Since the image of the specimen is blurred, we won't be able to acquire as much information about the structure of the specimen as we would with a corrected microscope. What's more, the information we can acquire will be harder to get, simply because it is harder to work with a blurred image of a specimen than a sharp one. There are very good pragmatic reasons for preferring particular types of instruments, and particular types of instrumental models. However, I cannot see how this poses any challenge to a traditional realist. All parties to this debate accept that instruments produce a variety of artefacts; that we need to keep track of the particular ways in which instruments transform information; and that sometimes we need to "correct" our instruments, i.e. alter the transforming processes so that they are more practically useful. Everybody who has experience with an instrument like a microscope will be painfully aware of these points, no matter what their philosophical commitments. If instrumental perspectivism is understood as involving the kind of "distortion" we find in spherical aberration, instrumental perspectivism is of little significance.

It is worth noting that Giere himself seems to slip from instrumental to theoretical perspectivism. We can see this in his discussion of two potential problems for the view, and his tactic for defusing these problems. The first potential problem discussed by Giere is that since instrumental perspectivism rests on a comparison to the human visual system, it overlooks the fact that the primary motivation for using instruments is to make scientific observation more reliable and objective than perception (Giere 2006: 41). We can use instruments to draw conclusions about matters that would be unclear or inaccessible to perception alone; and we often allow instruments to overrule conclusions drawn on the basis of perception, as when we might use a slow-motion film to determine the winner of a race. Instruments are useful precisely because they allow us to escape perspectivity. As Humphreys (2004: 10) puts it, "instruments trump human senses here, there, and almost everywhere."

Giere responds that this use of instruments is entirely compatible with instrumental perspectivism. A new instrument provides either a new perspective, as when a gamma-ray telescope reveals parts of the world inaccessible to perception, or perhaps a tool for making current perspectives more reliable, as when we use slow-

motion film to enhance the reliability of the visual perspective. Perspectives can be improved, and some perspectives are more reliable than others. On this point, I think Giere is partly right: there is no reason to think that the greater reliability afforded by some scientific instruments *in itself* undermines instrumental perspectivism, because the instrumental perspectivist can hold that some instrumental perspectives are superior to others. To say that perspective X is in some way better than perspective Y is not to say that perspective X provides an objective, non-perspectival “view from nowhere”. Of course, this raises the question of what the grounds are for judging that one instrumental perspective is more reliable than another: here we must appeal to the instrumental models and to our background theories of the systems to which the instruments are sensitive:

On what basis do I make all these claims about colour vision and scientific instruments? The short answer is, on the basis of various scientific theories; theories of colour vision, of electromagnetism, even quantum theory. But this answer immediately raises a further question: Are not scientific theories to be understood in an objectivist framework? If so, I remain an objectivist at the level of theory even if theory itself shows us that observation is perspectival. My reply is that theoretical claims are also perspectival. (Giere 2006: 59)

Giere deals with a potential objection to instrumental perspectivism by moving to theoretical perspectivism. His claims about instruments are made on the basis of various theories and models relevant to those instruments. But he says that this is no challenge to perspectivism, since he gives a perspectival account of theories and models as well.

The second problem, briefly touched on by Giere (2006: 49), is that even if instruments are perspectival, this is of little significance since scientists draw conclusions that go beyond the immediate data provided by instruments. There are, of course, epistemological problems with ampliative inferences, and such problems have long been discussed in philosophy of science. However, the perspectivist does not object to ampliative inference *per se*. Thus, the perspectivity of the data, or the partiality and opacity of the data, does not entail that any beliefs formed on the basis

of the data are similarly perspectival. Giere's response is that these ampliative inferences are made "only by moving to a broader theoretical perspective" (2006: 49). The fact that we draw conclusions beyond the data does nothing to show that instrumental perspectivism is false, because the theories and models on which such conclusions are based are also perspectival. Again, Giere attempts to save instrumental perspectivism by moving to a theoretical perspectivism.

The dependence of instrumental perspectivism on theoretical perspectivism may not in itself seem to be a problem; after all, I have said nothing to show that a perspectival account of the instrumental models is implausible. So far, my aim has simply been to make clear the relation between instrumental perspectivism and theoretical perspectivism. The former is dependent on latter – but perhaps the latter is true. Indeed, later chapters of this thesis will argue that theoretical perspectivism is true, in which case we get instrumental perspectivism for free. So why did I describe this dependence as a Pyrrhic victory for the instrumental perspectivist? Because it has two negative consequences.

First – and this is a minor point but still, I think, worth noting – it damages the overall structure of Giere's argument. As I noted in the introduction, in developing the case for theoretical perspectivism, Giere presents instrumental perspectivism as the less controversial thesis: he first argues for perspectivism about colour vision, then extends this argument to instrumental perspectivism, and finally, and "more controversial still", to theoretical perspectivism (Giere 2006: 14); see also an earlier work where Giere argues for perspectivism in observation and then claims that "the extension of perspectivalism to the level of scientific theory is more problematic" (Giere 1999: 81). If I am right, Giere has things backwards. Instrumental perspectivism rests on theoretical perspectivism.

Second, more importantly, the philosophical significance of instrumental perspectivism seems to be lost. Presumably, one would attempt to defend instrumental perspectivism specifically only if there is a reason for thinking of instruments as perspectival that does not apply to the many other entities and processes described by scientific models. So should we view instruments as being perspectival in some way that the host of other targets of scientific models are not? I

think this would be a mistake. An instrument is a nomological machine understood in terms of the instrumental model. The instrumental model tells us how the instrument operates and how it connects to the world. If such models are perspectival, then the claims we make about the world on the basis of our instruments will presumably be perspectival also. Similarly, if models of the Sun are perspectival, then the claims we make about the world on the basis of observed solar phenomena will presumably be perspectival also. The perspectivist doesn't take the Sun to be a perspective, and she shouldn't take instruments to be perspectives either. Attempting to apply the notion of perspective to the instrument itself adds little to our understanding of instruments and in fact creates puzzles, as I shall now outline.

2.5. Three puzzles about instrumental perspectives

2.5.1. The proliferation of perspectives

In Section 2.2.1, I argued that instruments are types of nomological machines to which we apply an instrumental model that links the instrument's output to processes in the world. But there are, of course, many other nomological machines: in particular, there are many nomological machines that arise in nature, and we may have models of these natural nomological machines which we are able to apply in order to draw conclusions about the world. Consider tree rings. Tree rings are an important source of evidence for past climate conditions, because tree growth is sensitive to environmental conditions and information about tree growth is preserved in the tree rings. Tree rings are the terminus of a well-ordered, well-understood causal process. This raises the question: what exactly is the distinction between scientific instruments and such nomological machines found in nature? The difference seems to me to lie in the fact that instruments are human artefacts and we can manipulate their parts relatively easily. Instruments are not metaphysically special and do not necessarily have any special epistemological status. If this is right, then it follows that if instruments are perspectival then so are all other complex, well-ordered causal processes.

Recall Giere's summary of instrumental perspectivism: "Observation does not simply reveal the intensity and distribution of gamma rays coming from the centre of the

Milky Way, it reveals the intensity and distribution of gamma rays *as indicated by COMPTEL or OSSE or...*" (2006: 48). Similarly, our conclusions about past climate conditions are drawn from tree rings; observation of the tree rings reveals the properties of past climates *as indicated by the tree rings*. Note that the information provided by tree rings is both partial, since trees are sensitive to only some features of the environment, and opaque, since the properties of tree rings are dependent on the growth and structure of the tree. I can't see any important epistemic distinction between the use of an instrument like COMPTEL by astronomers and the use of trees by dendroclimatologists. Should we say that trees provide "perspectives" on the climates of the past?

Perhaps Giere would accept the massive proliferation of perspectives that this position entails, that any nomological machine counts as a perspective. Indeed, Giere explicitly treats instrumental perspectives as being determined by human purposes; Giere (2006: 93) writes that instruments "are designed to interact selectively with the world in ways determined by human purposes." This is of course in line with Giere's broader perspectivism, which aims to put human purposes front and centre: his perspectival account of theories and models is built on an agent-based conception of scientific representation (Giere 2006: 60; Giere 2010). So contrary to what I have suggested, the instrumental perspectivist might be happy to take tree rings or even the Sun to be perspectives. After all, my own argument rests on the claim that instrument can be used and understood only with an instrumental model. Tree rings then can similarly become "instrumental perspectives" when we apply a particular kind of model to them to draw conclusions about past climate conditions. Along similar lines, the instrumental perspectivist can hold that it is our purposes that distinguish an instrumental perspective from the target system revealed by the perspective. There is no fixed line to be drawn; it simply depends on the research project in which the instrument is being used. The output of COMPTEL provides a perspective on nucleosynthesis, as opposed to the decay of aluminium-26, just when it is the former that we are studying.

But before the perspectivist bites the bullet on this one, it is worth emphasizing the damage that this bullet can inflict. It entails that any time we use a well-ordered system in the world to draw a conclusion about something else in the world, the

former is a perspective on the latter. It entails that the perspective of an instrument can be defined only relative to a particular scientist at a particular time. After all, two scientists in the same research group might use an instrument for slightly different purposes; and a single scientist might be interested in one process at one time, say aluminium-26 decay, and then a different but connected process at a later time, say nucleosynthesis. At this point, the question arises what work exactly the concept of “perspective” is supposed to be doing. What is the epistemological significance of the claim that any well-ordered system is a perspective, just when we use it to draw inferences about other systems? Once we understand the causal processes involved in a given nomological machine, processes that can be described using the same models by both realist and perspectivist alike, it seems to me that little is gained by describing it as a perspective. I suspect that the inclination is to say: the point of calling all these things “perspectives” is that they are all subject to limitations and distortions, in the sense I have described in the previous section, and so our access to the world is always from a particular point of view, and we cannot “detach” claims about the world from that point of view. Fair enough, but this point is already accommodated by theoretical perspectivism, by a perspectival account of the models that we apply when drawing inferences from phenomena – and recall that the instrumental perspectivist needs a perspectival account of these models in order to make sense of the notion that instruments distort.

2.5.2. The indeterminacy of perspectives

Visual perspectives are well-defined. A visual perspective shows us how something looks, or how something would appear to the eye, from a particular vantage point. There is nothing mysterious in supposing that an instrument occupies a specific visual perspective: we just need to describe what spatial location it occupies and which way it is pointing. The perspective of the instrument in the instrumental perspectivist’s sense is much more difficult to define; it seems to depend on how scientists are using the instrument and what their general goals are. Consider the title of one of the subsections of Giere’s chapter on instrumental perspectivism: “The Milky Way in Gamma Ray Perspectives” (Giere 2006: 45). Why suppose that we have a gamma-ray perspective on the Milky Way, rather than a perspective on gamma rays from the Milky Way? What’s the difference? The problem arises from

the fact that the behaviour of stars in the Milky Way, and the gamma rays produced by the stars, and the processing of the instrument that detects gamma rays, are all causal mechanisms involved in the production of the instrument's output. The "perspective" seems to include whichever part of the causal chain we are taking for granted when investigating some earlier part of the causal chain.

Another way to make this point is to ask what a given perspective is a perspective on, i.e. what the target of the perspective? Take an image created by COMPTEL, such as an image constructed from data of gamma rays at 1.8 MeV. Perhaps it represents the Milky Way galaxy. Perhaps it represents the decay of aluminium-26, since gamma rays of 1.8MeV are produced by the decay of this element. Perhaps it represents sites of nucleosynthesis in the cores of massive stars, as this is where aluminium-26 is primarily produced and its concentrations in the galaxy are greatest where nucleosynthesis is still occurring. Perhaps it represents sites of star formation, as massive stars tend to be closer to stellar nurseries. Perhaps it represents the regions of gas and dust throughout the Milky Way that obscure gamma rays, much as a photo can represent a silhouette that blocks out background light. Perhaps it represents the scintillations taking place inside the machine that are measured by the photomultiplier tubes. The image could represent all these things. It could be a perspective on all these things.

What is required to solve this problem is some principled way of drawing a line between *an instrumental perspective* and *the systems in the world the instrumental perspective reveals*. We need to be able to say: *this* set of causal processes constitutes the perspective, and *this* set of causal processes are what is revealed by the perspective. I don't know what grounds there could be for drawing such a line, or what the epistemological significance of the distinction would be. Certainly, nothing in our best scientific theories and models provides any guidance for drawing this line.

If anything, some analyses of scientific disputes suggest that the development of science often depends on this line being vague. I am thinking here of Pinch's (1985) point about how observation reports in science can have different "degrees of externality". Whenever a scientist uses an instrument, she faces the question of how the output should be described. To take Pinch's example, one type of solar neutrino

detector consists of a tank of perchloroethylene, which contains a chlorine isotope with which neutrinos can interact in such a way as to produce radioactive argon. By counting the argon atoms, we can determine the solar neutrino flux. When a scientist uses this detector, she might report from these results that: (1) she observed splodges on a graph, or (2) she observed radioactive decay, or (3) she observed the decay or Ar^{37} isotopes, or (4) she observed solar neutrinos. As Pinch says: "These reports can be said to exhibit progressively increased externality ... in the sense that more and more of the observational situation must be encompassed in order to produce the report" (1985: 9). In the first case, the scientist needs to trust only her own perceptual system. In the second, she must trust that she has an accurate instrumental model of the instrument used to detect radioactive decay, which may just be a simple Geiger counter, for example. In the third, she must trust that she has an accurate instrumental model of the whole perchloroethylene tank and argon capture process, of which the previous instrument is a part. In the final case, she must assume that she has an accurate model of neutrinos and their interaction with the detector. With greater degrees of externality, the further the instrumental model is extended.

In deciding how to report an observation, scientists face a trade-off between security and profundity. A report of low externality pertains to instruments that are well-understood, and so is more likely to be correct and so to go unchallenged, but for precisely this reason it will be less interesting. A report of high externality is likely to be profound, but also risky. The higher the externality, the more opportunities there are for criticism. Pinch argues that many scientific disputes involve attempts to alter the externality of observation reports (1985: 15). One challenge to the original solar neutrino experiments targeted the model of the capture of neutrinos by chlorine. Had this challenge been sustained, the experiments would have supported reports of type (3) but not (4). Another challenge suggested that the argon chemistry had been misunderstood, and the argon isotopes will have remained trapped in the tank. Had this challenge been sustained, only reports of type (2) but not (3) could have been given. A great deal of debate in the sciences therefore rests on where to draw the line between an instrument and the instrument's target. Even before debate can begin, a scientist must decide whether she wants to make a profound but risky claim or a less interesting but more secure one. Clearly there can be no algorithm for this;

indeed, it will depend to some degree on the personality and values of the scientist in question.

2.5.3. Interacting perspectives

A third puzzle arises when we consider how different perspectives interact. An interesting case that provides a good illustration of this is found in Vertesi (2015), who describes the process of creating images of Mars. The Mars rovers Spirit and Opportunity are both equipped with panoramic cameras, or PanCams; each PanCam has thirteen optical filters. Often scientists will capture several photographs of the same scene, each with a different optical filter. By combining these photographs in different ways, different aspects of the scene are made visible, revealing the various chemical and mineralogical properties of the objects. For the instrumental perspectivist the result is the surface of Mars from the perspective of the PanCam.

However, as Vertesi discusses (2015: 53-72), all images received from the PanCam undergo a calibration process where they are corrected to account for the local conditions on Mars. This is done using a “caltarget”, an object painted with red, green, blue, and yellow sections, and three shades of grey. Since the values for the colours on the caltarget are precisely known, images of it allow scientists to determine exactly how local conditions affect the colours in other images. With the aid of computer analysis, they develop equations that correct for light pollution, scattering, dust, and other conditions in a particular set of photos. Part of the calibration procedure crucially involves humans, since before the computer analysis can occur, humans must highlight different parts of the caltarget image: “the calibrator must identify the different colored zones on the caltarget for the computer, so that the computer can then calculate how much each individual image can be adjusted” (Vertesi 2015: 59). The calibrator must select regions of pixels and tag them as corresponding to specific colour zones of the caltarget. Computers are not yet nearly as efficient at such visual judgements as humans are, so humans are crucial for this stage of image processing.

What is the perspective to which the final output is indexed? If we say that we have an image of Mars from the perspective of the PanCam, we are ignoring all the other processes involved in image construction, such as the crucial role of computer analysis and human judgement. The humans employ their visual colour perspectives when identifying zones on the caltarget images. These human perspectives are essential for generating the data that scientists use to draw conclusions about Mars. In fact, many of the scientists involved in the Mars mission are interested not in the corrected images but in the calibration procedure itself: as Vertesi notes, atmospheric scientists and soil scientists are often interested in properties of the atmospheric dust, and so “They therefore use the output from the calibration procedure to get dust information and would rather see the dust than the image it obscures” (2015: 78). This output surely can’t be indexed simply to the perspective of the PanCam. Instead, perhaps the output is a product of the PanCam perspective plus the human colour perspective. Or do these combine into a new perspective entirely? Nothing in Giere’s definition of perspectivism or his subsequent discussion rules out combining perspectives to produce novel data. But how do we decide whether we have two perspectives working in tandem, or a wholly new perspective constructed out of two perspectives? It is hard to see how we could go about answering this question – or what difference the answer would make if we found it.

Unless the instrumental perspectivist finds some principled way to answer these problems, then for any instrument, and any output of that instrument, we cannot specify what exactly the perspective of the instrument is. Instrumental perspectivism is ill-defined. It is much more straightforward to simply think of the instrument as a causal process, as outlined in the instrumental model. To be clear, my goal in the previous sections has not so much been to show that instrumental perspectivism is incorrect, but that it is unmotivated. If the theoretical perspectivist is right, then models are perspectival, and this of course includes the instrumental models, but in this respect instruments are no different from any other system we can model. To attempt to apply the notion of perspective to the instruments themselves raises various puzzles. For both realist and theoretical perspectivist, there is a far simpler way of interpreting the situation: we have models of various phenomena, including instruments, and we use those models in addition to our observations to draw

conclusions about various systems in the world. The perspectivist then adds a perspectival analysis of all these models.

The direct argument for instrumental perspectivism fails. But there may still be resources to defend a kind of instrumental perspectivism, which avoids the problems outlined here, in the form of an indirect argument that appeals to extended cognition. I turn to this in the next chapter.

Chapter 3

Extended Perception

3.1. Do instruments extend perception?

In the previous chapter, I examined Giere's approach to instrumental perspectivism and found that it faces some serious problems. However, as I pointed out in the introduction to that chapter, there is an alternative way to think about instrumental perspectivism, which allows the instrumental perspectivist to avoid these problems. The "indirect argument" for instrumental perspectivism appeals to the resources of the extended mind thesis and holds that instruments are literally part of our perceptual processes. If Giere is right that perception is perspectival, and here we are granting this for the sake of argument, then instruments must also be perspectival as they are simply another kind of perception.

To be clear, this argument is not endorsed by Giere. However, Giere does make some suggestions in this direction. For example, in an article on the scientific realism debate, he appeals to work in distributed cognition to attack van Fraassen's constructive empiricism (see Giere 2005: 153). Van Fraassen holds that the aim of science is truth about the observable, where the observable are those things that can be detected by the unaided human senses. Distributed cognition challenges van Fraassen's distinction between observable and unobservable. A human is a cognitive system, but on the distributed cognition approach, so also is a human plus instrument. On van Fraassen's view, we treat the instrument merely as input to the human system. But treating the human plus instrument as one unified system, it no longer makes sense to privilege the output of human senses. A single human alone cannot detect air pressure; a human plus air pump can.

This chapter will examine the extended mind hypothesis as an alternative route to instrumental perspectivism. One thing that may be quite appealing about this approach is that it gives us instrumental perspectivism "for free", as it were. Giere has already argued for the perspectivity of perception, and there are a variety of independent reasons, beyond the concerns that motivate perspectivism in general, to favour the extended mind. With this in place, all we need to do is show that

instruments extend perceptual processes, and we have established instrumental perspectivism.

There is clearly a close connection between instrumentation and perception. In particular, instrumentation and perception share one very important feature: they are both means for detecting properties in the world by establishing a special kind of causal connection with those properties. Indeed, we might view our perceptual systems as being types of instruments. Imagine a race of blind aliens who use human reports of colours as data against which they test their scientific theories. Humans are subjected to a certain type of input, and they produce an output, phrases like “the ball is red”, “the cube is orange”, “*a* is darker than *b*”, “*c* is bluer than *d*”, etc. Just as we use instruments such as gamma ray telescopes, and link particular data produced by those telescopes to specific wavelengths of gamma rays, so the aliens could link our colour reports to specific wavelengths of electromagnetic radiation.

So perhaps perception just is an instrument. But are all, or a significant number, of the instruments beyond the body also perceptual? The application of the extended mind to scientific instruments is defended by Toon (2014). Toon discusses the debate between realists and constructive empiricists concerning the distinction between observables and unobservables. Realists argue that instruments such as microscopes allow us to observe entities such as bacteria; for realists, such instruments are, as Van Fraassen (2008: 96) puts it, “windows on the invisible world”. In contrast, constructive empiricists hold that the observable is simply that which is revealed by the senses, and instruments are “engines of creation” that produce new images that must be accommodated by scientific theories, but there is no need to assume that these images correspond to any real objects. Toon, in an argument similar to Giere’s appeal to distributed cognition, suggests that the extended mind thesis challenges the constructive empiricist position. It is not that instruments simply augment the senses, as the traditional realist would argue; rather, instruments are literally part of the perceptual process.

As I will discuss later, there are many different types of extended mind hypothesis, which could allow different ways to defend the claim that instruments extend

perception. For now I will focus on Toon's argument which appeals to the parity principle of Clark and Chalmers (1998):

If, as we confront some task, a part of the world functions as a process which, were it done in the head, we would have no hesitation in recognizing as part of the cognitive process, then that part of the world is (so we claim) part of the cognitive process. (Clark and Chalmers 1998: 8)

One way to think about the parity principle is to imagine an intelligent Martian with a mechanism in its head that functions just like the instrument used by the scientist. If we would judge this to be part of a perceptual process in the Martian, we should also conclude it forms part of a perceptual process for the scientist. The argument for the claim some instrument is part of a perceptual process is then straightforward: if the instrument were inside the head of a Martian, it would count as part of the Martian's perceptual process; hence by the parity principle, the instrument is part of the scientist's perceptual process (Toon 2014: 415). So imagine a Martian with a mechanism inside its head that can detect what is detected by the human eye plus the optical microscope. In the case of a single-lens microscope, we can just imagine a Martian with a visual system that is essentially the same as the human visual system, except the lenses of its eyes have a far smaller diameter. Surely we would say that this Martian can see paramecia. A similar story can be told for almost any other scientific instrument – or so it seems, at least. For any instrument, we can imagine a similar mechanism inside the head of an intelligent Martian.

Before proceeding it is worth noting a potential snag in this argument. The argument that instruments extend perception rests on the extended mind hypothesis, which is itself arguably a scientific hypothesis. As many critics of constructive empiricism have noted, it's questionable whether constructive empiricists can consistently rely on the sciences to define what is perceivable, what is revealed by the senses, given the agnostic stance towards theories that constructive empiricism recommends (for discussion of this see Muller and van Fraassen 2008). This is an interesting problem, though I don't think it arises for the perspectivist. There are no obvious problems with a perspectivist relying on our best theories to define the boundaries of perception, since the perspectivist accepts that our best theories are true. Here then we will

assume that the extended mind thesis is correct, and that it is acceptable for the perspectivist to appeal to this in defining perception. Does the extended mind thesis apply to scientific instruments? The rest of this section is structured as follows. First, I develop an alternative way of thinking about scientific instruments that still respects the insights of the extended mind. Second, I criticize the use of the parity principle to argue that instruments extend perception.

3.2. An alternative view of instruments

Many extended mind arguments involve the extension of capacities that we already have. Toon cites the famous example of Otto and Inga given by Clark and Chalmers (1998). Both Otto and Inga want to visit the Museum of Modern Art. Inga is a normal adult, and she simply remembers that the Museum is located on 53rd street. Otto has Alzheimer's and carries a notebook with him in which he records important information. He looks up the location of the Museum and sees that it's on 53rd street. Clark and Chalmers argue that the notebook is literally part of Otto's memory, and that whatever is written in Otto's notebook forms part of his set of beliefs.

Another example is the use of writing while working out problems. This simple activity can promote creative thinking: writing lots of notes while working on an essay can prompt new ideas or help you to see unexpected connections between old ideas. Clark compares this to turbocharging (2008: 131). A turbocharger uses the exhaust flow output of an engine to power an air pump that compresses the air flowing into the engine, thus improving the engine's performance. The turbocharger counts as a part of the vehicle's power-generating mechanism. In the same way, the brain sometimes produces outputs, such as written words, that it then recycles as inputs to improve the cognitive process. When you write, you are not just recording your thoughts, you are turbocharging them; hence note-taking becomes part of the cognitive activity.

In these circumstances tools beyond the head are used to extend and augment the cognitive capacities we already have. But if scientific instruments extend perception, this produces new capacities; the instruments give us, quite literally, new senses. This may not in itself be a problem. For now, I only wish to note that there is a

significant difference between perceptual augmentations such as eyeglasses and the instruments of science.

Actually, not all scientific instruments differ in this way. It seems to me that, if the extended mind thesis is right, it would be the case that the optical microscope extends *visual perception*, so that we literally see with a microscope. The optical microscope shows us an object in visible light, using a system of lenses to magnify the object. This is very similar to the mechanism of the human eye, which uses a lens to refract visible light. Indeed, the simplest microscopes such as those of Van Leeuwenhoek were just single lenses. A spherical bead of glass acts as a magnifier: the smaller the diameter, the greater the magnification. Van Leeuwenhoek's surviving lenses achieve a magnification of up to 275x; some of his lenses possibly achieved magnification of up to 500x, an exceptional magnification for the time (Bradbury 1967: 73). By comparison, modern compound microscopes can magnify up to 800x with lenses alone, without the use of techniques such as oil immersion. Such visual systems already exist in some small animals. A small fish may see a paramecium that is invisible to the human eye. In this case the parity principle is straightforwardly applied. We can certainly imagine a Martian with an optical microscope in its head – or just a Martian with lenses in its eyes of a smaller diameter than are found in humans. We would say that the Martian sees the paramecium. So, by the parity principle we should say that the optical microscope literally extends vision. The question is just whether the optical microscope meets enough other conditions for a perceptual process, e.g. "glue and trust" conditions. I think there are plausible responses to this point (cf. Toon 2014: 422); in any case I put it aside for now. Let's grant that the optical microscope extends vision and thus inherits the perspectivity of the visual process. What about other instruments?

The resolving power of optical microscopes is limited by diffraction. We can achieve greater resolving power with different mechanisms. One such mechanism is found in atomic force microscopy. In brief, this works essentially as follows. The microscope contains a cantilever, a long rod with a sharp tip at the end; the point of the tip consists of just a few atoms. The cantilever is moved across the sample, and the forces between the tip and the surface deflect the cantilever up and down, depending on the topography of the sample. A laser beam is pointed at the top of the

cantilever, and the beam reflects from the top of the cantilever into a detector. As the cantilever moves up and down, it deflects the laser beam; and by tracking how the laser is deflected, a computer can build up an image of the topography of the sample.

Obviously, this is nothing like vision. Perhaps it is more analogous to touch? Not really. Remember that we are on an atomic scale here. The repulsive intermolecular forces push the tip away: these forces, and how they interact with the cantilever, will differ depending on the sample. Any inferences about the measurements of the topography of the sample are therefore based on physical theory concerning the intermolecular forces of the sample and how they will affect the cantilever tip. So what exactly is it that is being extended here? The atomic force microscope produces an image. This is visually perceived. We should not be misled by this into supposing that atomic force microscopy extends sight. The information produced by an atomic force microscope could be communicated in some other way – via sound or touch. In fact, its mechanism is nothing like any of the senses. The same is of course of many other scientific instruments. If we are to say that the atomic force microscope extends perception, it does not do so by extending any perceptual capacity we already have, but instead creates an entirely new perceptual capacity.

Toon is aware of this point but argues that it does not challenge to extended perception argument: “the extended perception argument does not rely on the claim that the aliens perceive objects using the same physical processes, or even with the same fine-grained perceptual psychology, as we do” (2014: 421). Consider again the Martian with an atomic force microscope in its head. We don’t need to claim that this is a form of sight. All that matters is that it counts as some sort of perceptual process. This is surely plausible: after all, we already recognize that different creatures may have perceptual processes very different from anything we find in humans, such as echolocation in bats, magnetoception in birds, and electroception in sharks. So the fact that most scientific instruments would have to be taken as producing new perceptual capacities rather than simply extending old ones does not in itself refute the extended perception argument. But I do think that this difference between the optical microscope on the one hand, and the atomic force microscope on the other, points to a serious problem for the argument.

The parity principle makes sense for optical microscopes. It's notable that even van Fraassen is rather nonchalant about the view that we observe through optical microscopes (2001: 162-163): for him what's important is that it makes sense to draw a line, even if the line should not be drawn exactly where he draws it. I will argue that we should draw a line. As I see it, the significant difference between the optical microscope and the atomic force microscope is that when we use an atomic force microscope to detect some X, the detection of X is necessarily *indirect*. The atomic force microscope produces a literal image on a computer screen, and we see X only via seeing the image. The information about X could of course be delivered in some other way, as the microscope could produce sounds or construct an object that could be touched; but our access to X would always be indirect in the same sense. We access X only via an artefact produced by the microscope.

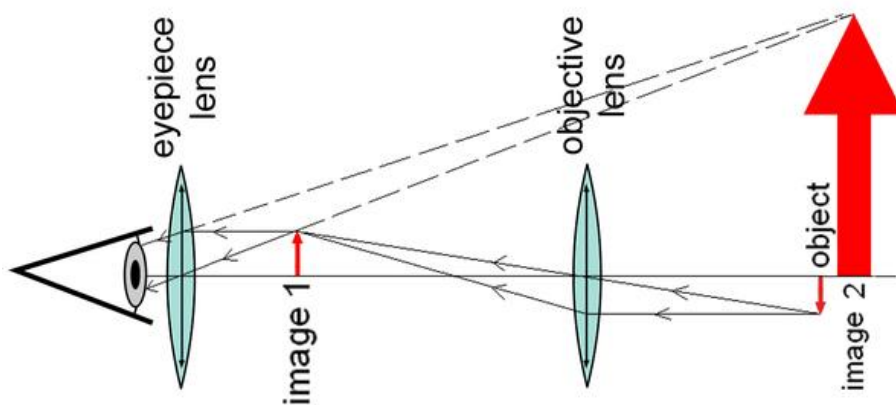
In contrast, when we look at X through an optical microscope, we are not looking at an artefact that represents X. If I use a lens to view a paramecium, I can see that the lens is transparent, that there is not a literal image of a paramecium inside the lens. (We may talk about lenses creating real images and virtual images, but these technical terms refer respectively to where light rays converge and where light rays appear to have converged; real images and virtual images are not material objects.) Van Fraassen (2008: 103-105) provides a catalogue of images. *Graven images* are images that are material objects: paintings, photographs, sculptures, films. *Private images* are purely subjective experiences such as dreams, hallucinations, and after-images. Finally there are *public hallucinations*, which are images that are not subjective but are also not genuine objects: rainbows, shadows, reflections, mirages. Some public hallucinations are images of real things (the reflection of the tree in the lake); others are not images of anything (rainbows are not images of anything). Van Fraassen treats the images produced by optical microscopes as public hallucinations. The question is, should take such images as images of real things?

The category of public hallucinations is questionable, but an evaluation of this category is not exactly my point. Let's say we hold that this category should be rejected. Then it seems we have no choice but to treat optical microscopes as "windows on the invisible world." There is nothing else they could be. They do not

create graven images. There is not literally an image, a material object, inside the lens of an optical microscope; no more than there is a literal image inside the lens of a magnifying glass. In this respect a simple microscope is no different in kind from a powerful magnifying glass or even a window. We simply see through these things; they do not produce images. In these cases, our perception of the object is direct. As Maxwell (1962: 7) said, there is a continuous series from “looking through a window pane, looking through glasses, looking through binoculars, looking through a low power microscope, looking through a high power microscope, etc.” – this is exactly right provided we are considering only optical microscopes, no matter what their magnification.

Many scientific instruments produce only graven images, which may or may not depict real things: the electron microscope, the atomic force microscope. Other scientific instruments do not produce images at all, or at least their outputs are not usually interpreted as images: the white streaks produced by a cloud chamber do not represent charged particles, but according to theory are simply produced by charged particles much as a jet plane produces a trail in the sky. Now let me state my claim: In Van Fraassen’s terminology, only instruments that create public hallucinations are plausible candidates for extending perception. Instruments that produce things, even if these things are graven images, as in the case of the atomic force microscope, are not. The line between instruments that extend perception and those that do not is where the only access to the information provided by the instrument is via a graven image. In my view the atomic force microscope has more in common with a cloud chamber than with the optical microscope. The atomic force microscope produces an image that we take to represent some unobservable structure. The cloud chamber produces phenomena that we could take as representational, but we do not. Now peer through an optical microscope. There is no image, there is no artefact about which we can ask: what, if anything, does it represent? This is a question that may be asked about the output of an atomic force microscope. This is a significant difference, completely lost if we view scientific instruments in general as extensions of perception.

It’s worth noting that this difference is presupposed by scientists themselves. Shown below is the instrumental model of a compound optical microscope:



This model includes the human eye at the left. The model must depict a “receiver” at the left; without this, it would not properly show how the optical microscope functions. This receiver need not be a human eye, but it can be. Contrast this with the instrumental model of an atomic force microscope, where no parts could be replaced with a human sensory organ.

Perhaps the perspectivist could grant that there is an important difference here, but still insist that I have not really dealt with the extended perception argument. After all, everybody recognizes that there are differences between different perceptual processes. Sight is certainly not the same as touch. So, if we must recognize significant differences between perception with an optical microscope and perception with an atomic force microscope, this is no matter. I grant that so far, I have only presented an alternative view; I have not yet dealt with the appeal to the parity principle. Let’s turn to this.

3.3. *Why the appeal to the parity principle fails*

Imagine a Martian with an atomic force microscope in its head: surely, despite what has been said so far, we should say that this is part of the Martian’s perceptual process. Hence, by the parity principle, the atomic force microscope does extend perception! I am not convinced. The thought experiment with the Martian is seriously misleading in two ways.

First, what we are tempted to imagine when we imagine the atomic force Martian is that the output of the atomic force mechanism in the Martian’s head is not an artefact

that is then made available for the Martian's examination, but is simply immediate sensory data. But this, of course, is just to ignore the difference that I have outlined above. When we adjust the thought experiment to account for this difference, we do not have the intuitive judgement that the Martian is perceiving anything.

Consider photography. A photograph creates an image of a scene that is highly realist. Ideally a photograph looks exactly like the original scene looks, if you are viewing the scene from the same vantage point as the camera. When we look at a photograph, we know that it accurately represents what it depicts, provided it hasn't been retouched. So this is a case where questions about realism and antirealism do not arise; we all agree that photography can augment our perceptual capacities since we could in principle perceive the objects depicted in a photograph without the aid of the photograph. Yet Walton's (1984) claim that photographs are transparent, that we literally see through a photograph to the object it depicts, at least initially seems bizarre. Why? What photographs lack is a certain kind of directness. When we see a photograph of an object, we see the object only via seeing another, completely different object.

Does the extended mind thesis support Walton? I don't think so. Imagine a Martian that has in its head a mechanism that functions just like a camera. This mechanism produces an image of whatever the Martian is looking at – a literal image, a graven image, that is stored inside its head in a file of images that it can flick through when it wants to revisit them. One day the Martian witnesses a crime, it creates several images of the crime, which later it revisits to answer questions about the crime. Intuitively this is a cognitive process. But is it perception? Surely it would more naturally be classed as a kind of memory – literally a photographic memory! What this thought experiment establishes, if it establishes anything, is that there are circumstances where a file of photographs can count as an extended memory. *Perception* through a photograph seems no less bizarre when it's going on in the head of a Martian than under normal circumstances. The moral for the atomic force microscope should be clear: if even a photograph, which we all agree accurately reveals the genuine structure of the world, does not extend perception, then an image produced by an atomic force mechanism surely doesn't either.

I have claimed that what is missing in photographs and in images created by atomic force microscopes is a certain kind of directness, since the object is perceived only via perceiving another object. What exactly do I mean by *directness*? Unfortunately, I do not know how to define this for perception in general. To begin developing the notion I suggest that visual perception of an object requires the detection of light that has interacted with an object and/or its surroundings. I see a mountain by detecting light that has been reflected by the mountain. Light hits the mountain: some is absorbed, some is reflected and hits my retina. Perhaps we would say that an object is seen when it blocks light, as when we see a silhouette; here we detect the object by detecting the light around it. Either way, what hits our eyes is light that has interacted in some way with the object or its surroundings. This condition holds for objects seen through optical microscopes: viewing a paramecium through an optical microscope, the light that hits my eye is light that was reflected off or transmitted through the paramecium. Hence we have a direct connection to the paramecium. However, looking at a graven image of a paramecium created by a photograph or an atomic force microscope, I am not detecting light that was reflected off or transmitted through the paramecium. This suggests a necessary condition for visual perception (though not a sufficient condition: light is transmitted through air, but we do not visually perceive the air).

Atomic force microscopes lack directness in another way also. We have noted that if atomic force microscopes extend perception, they do so by creating a new sense. What's odd about this new sense is that it is entirely parasitic on old senses. We perceive the nanoscale object via perceiving another object, and that other object is perceived through sight.

It may be pointed that senses are already dependent on each other. Our normal senses interact, sometimes very strongly, so that one sense relies on another. Taste is strongly influenced by smell. But here we are dealing with interactions. We can still taste without smell, and taste is still a different sense to smell. Importantly, we don't taste via the nose. Taste and smell are phenomenologically very different experiences. However, looking at an image created by an atomic force microscope is phenomenologically like looking at any other image. It feels like seeing. Indeed, there is surely no question that this is what we are doing. Yet if the atomic force

microscope extends perception, the perceptual process is not a form of *seeing*. The oddness of this situation is brought out very clearly if we imagine looking at two graven images that are qualitatively identical, where one is produced by an atomic force microscope, and the other is a pattern of random noise that just by chance created the same image. To look at the former is to engage in a wholly new form of perception, to look at the latter is not. It may be objected: hallucinating a tree has the same phenomenology as seeing a tree, but they are still completely different processes. This is true. But the analogy fails, because a hallucination is a purely internal process; it involves an experience like a perception but without an external stimulus. In the case of the two graven images, however, it is agreed by all that in both cases we are perceiving the world.

The second problem with the Martian thought experiment is that given how instruments such as atomic force microscopes function, it is completely implausible to suppose that they could ever be part of the head of any organism. Of course, philosophers often appeal to thought experiments, sometimes quite outlandish. The real problem in this case is that it's not at all clear what we are being asked to imagine. We are tempted by the image of a Martian with atomic force microscopes in place of its eyes; as it wanders around the world it points its atomic force microscopes at things and then receives information about them. Intuitively, this is a perceptual process. But of course, this is simply a fantasy.

You do not point an atomic force microscope at anything; you place a sample under the cantilever. Such samples must be carefully prepared: first the sample must be cut to an appropriate size; then it must be fixed to an appropriate substrate to prevent any movement, which would otherwise interfere with the measurements made by the cantilever; chemicals may need to be applied to make certain features more salient; etc. The same is the case for other types of sophisticated microscopy. When using such instruments we first must create an artefact that can interact with the microscope in the right kind of way to produce a usable image. This is where the "window on the invisible world" metaphor breaks down. Even if there were a way of getting an atomic force microscope inside the head, such a creature could not use it in its interactions with the world in anything like the way any other perceptual process is used.

It would be one thing to say that the atomic force microscope extends capacities that we already have. I have already argued however that this is not how it should be viewed: if it extends perception, it does so by creating new perceptual capacities. This raises a crucial question: what exactly is a *perceptual* process? We might say that a process counts as perceptual just in case it allows us to detect features of the world; an instrument can therefore extend perception just when it allows us to detect features that would otherwise be undetectable. But notice that this is, for a realist, also true of theories. Modern astrophysical theory allows me to learn about, say, the properties and distribution of dark energy, so if there is dark energy our minds are detecting it, even if it is otherwise totally undetectable with our best instruments. Do I perceive dark energy? If so, the very notion of perception as distinct from cognition collapses. Perhaps this is a consequence some people would be willing to embrace. But it spells serious trouble for the perspectivist argument, where the claim is that perception is perspectival and this can be extended in a straightforward manner to scientific instruments. If we can no longer even distinguish perception from cognition, this argument will have to be reformulated. I assume then that perception requires more than just detection. Note that this is tangential to extended mind argument: even if you reject the extended mind, perception requires more than detection plus processes occurring in the head.

When we imagine a Martian with an optical microscope in its head, we are imagining an organism that is very much like organisms that we are familiar with. For this thought experiment we can simply imagine an organism with lenses of a smaller diameter in its eyes. But when we try to imagine a Martian with an atomic force microscope, or with some other mechanism that magnifies as much as an atomic force microscope, what would this organism be like? Would we still suppose that there are any *perceptual* processes taking place? The previous example of the Martian with the camera in its head, which intuitively counts as a kind of memory rather than perception, should illustrate that we cannot rely on our intuitions about whether some mechanism is a perceptual process until we are able to state specifically how the mechanism functions within the organism. Perhaps it would be possible to provide a thought experiment where the mechanism is described in detail and where our intuitive judgement is that it is a perceptual process. But then it would also be possible to provide a thought experiment where the mechanism is

described in detail and our intuitive judgement is that it is not a perceptual process. I conclude that the parity principle does not support the claim that instruments such as atomic force microscopes extend perception.

3.4. Beyond the parity principle?

There are, of course, other arguments for extended perception that do not rely on the parity principle. Toon points out (2014: 416) two other arguments for the extended mind that appeal to the notion of close coupling (Clark 2010) and to phenomenological considerations (Clark 2003). These can be dealt with swiftly. Close coupling is not possible for sophisticated instruments such as atomic force microscopes. Perhaps it would be possible for a person to carry a simple optical microscope everywhere and to employ it enough in their dealings with the world that it counts as part of their cognitive system. Few instruments are so user-friendly. As for phenomenology, I have already commented on this. Looking at an image produced by an atomic force microscope is just like looking at an image; it does not have the different phenomenology that would be expected of a new perceptual process.

There is another option that is not discussed by Toon: enactivism as defended by e.g. Noë (2004). Enactivism proposes that perception does not involve the mind mirroring or representing the world but is instead a relational process arising from the interaction between the organism and its environment, from the organism exercising various skills and abilities in the environment. Perception is tied to action; we perceive the world in order to do various things in it. In particular, perception involves exercising our knowledge of “sensorimotor contingencies”, which are regularities in the way that sensations change as a result of bodily movements. As I move towards a black object, for instance, my visual sensations will change in regular, predictable ways: blackness will increasingly fill my visual field, and blackness will occlude more objects previously in my visual field. In general, I know that an object will appear larger as I approach it, I know how its apparent angles will change as I move around it, I know how it will feel if I pick it up and how much force to use when doing so, I know how moving will reveal things occluded by the object, etc. In Gibson’s

language, we see the environment as *affording* various possibilities for action (Gibson 1979).

Enactivism provides no comfort for the instrumental perspectivist. If anything, my claim about the fundamental difference between perceiving X and perceiving a graven image of X is especially well-suited to the enactivist view. On my desk is a cup and a phone. I know that as I move, different parts of the cup and phone will be revealed; that at times the cup will occlude the phone; that some parts will be in shadow and others in light. The cup and the phone both afford various possibilities for action, in that I could throw the cup and use it to smash the window. Now suppose I look at a photo of the same scene. As I move around the photo, the image of the cup doesn't occlude the image of the phone. I can't use the image of the cup to smash the window. The affordances of X and the affordances of an image of X are very different. The output of an atomic force microscope is merely an image; the sensorimotor contingencies are just as they are for a photograph or a painting. I conclude that there is not yet any good reason to think that instruments in general literally extend perception. Perhaps there are a few examples, such as optical microscopes, but none of the major theories of extended perception apply to all or even most scientific instruments.

3.5. Conclusion

I will now summarize the results of the previous two chapters. At first blush, instrumental perspectivism probably seems to be the most plausible form of perspectivism in the sciences. Giere treats it as the less controversial thesis, and as the natural extension of perceptual perspectivism. I believe that closer examination reveals that it is deeply unattractive.

We saw in the previous chapter that Giere defends instrumental perspectivism by appealing to two properties of instruments: partiality and opacity. Partiality is trivial and can clearly be accepted by a standard realist. Opacity is a much more intriguing concept but seems to conflate transformation of information with distortion of information. Like partiality, transformation is trivial: some sort of transforming process is required for an otherwise unobservable property to be made detectable. I take it

then that the key claim in the direct argument for instrumental perspectivism is that instruments in some sense distort. I then argued that an instrument can be said to distort only when there are inaccuracies in the instrumental model, and that from this it follows that instrumental perspectivism is dependent on theoretical perspectivism in a way that drains its philosophical significance. Furthermore, attempting to apply the metaphor of perspective to the instrument itself generates puzzles concerning how perspectives are to be specified.

The prospects are equally poor for defending instrumental perspectivism by appealing to the extended mind. Even if the extended mind hypothesis is plausible, and even if perception is perspectival – both of which are, of course, extremely contentious – there is no convincing reason to treat instruments in general as methods of extending perception. Only a few relatively simple instruments, such as some optical microscopes, are good candidates for this.

So, both approaches to instrumental perspectivism, the direct argument and the indirect argument, are seriously flawed. A key point that I have emphasized throughout these chapters is that an instrument is well-ordered causal process, and to draw any conclusions about the world when using an instrument, scientists must understand the instrument in terms of an instrument model that describes how this causal process works. I don't think this is especially controversial; indeed, even the perspectivist would agree with this. But now this opens the instrumental perspectivist up to a straightforward application of Chakravartty's *Escape from Perspective*. We learn non-perspectival facts about the world, using scientific instruments, when we have accurate instrumental models. At this point we can ask: What do we then gain by describing the instrument as a *perspective*? What problems does this solve? What does this bring to our understanding of the instrument or the properties it detects? Not much, as far as I can tell. It does however raise various difficulties. Instrumental perspectivism introduces new problems while providing few benefits. I conclude that philosophers attracted to perspectivism would do better to explore the prospects of theoretical perspectivism instead. That will be the task of the remainder of this thesis.

III

THEORETICAL PERSPECTIVISM

Chapter 4

Incompatible Models and Perspectival Truth

4.1. Multiplicity

In this second part, I will turn to what I have called “theoretical perspectivism”. Since theoretical perspectivism will occupy us for the rest of this thesis, henceforth I will refer to it simply as “perspectivism”.

As we have seen in the introduction and in the previous part, the underlying intuition that drives perspectivism is the idea that truth is revealed from multiple viewpoints. Call this the “multiplicity intuition”. There are, of course, different ways of elaborating on this intuition. On a traditional realist elaboration, we can say that theories are *selective*. Theories track mind-independent facts in the world, but naturally, no theory is expected to provide a complete account of the world. Our cognitive capacities are far too limited for that. Instead, we focus on particular features of interest, which will be different in different contexts; and we abstract from precise details to identify higher-level patterns, where again, the patterns of interest depend on the context. As a result, we develop a multitude of theories and models, a multitude of perspectives for understanding the world. This is all entirely compatible with realism as traditionally understood.

On an antirealist elaboration, by contrast, we can say that we get at the truth in specific domains via the means of falsehoods in other domains. For example, we construct theories that provide true accounts of the observable phenomena, though the underlying theoretical claims may be false. Newtonian mechanics provides a “perspective” on the observable phenomena: it is a tool for systematizing, predicting, and controlling that which impinges on our senses. General relativity provides an alternative perspective, one which supports more accurate predictions, but even so, there are still plenty of contexts in which we use the Newtonian perspective instead. As before, theories track mind-independent facts, but these are facts about the

observable world specifically. Stronger forms of antirealism may say that science does not get at the truth at all, that there is a multiplicity of useful falsehoods.

The point of this is that the multiplicity intuition is widely shared; it is something to be accounted for within any acceptable philosophical account of science. We must recognize the plurality of perspectives that science provides for understanding and interacting with the world. This might support the initial plausibility of perspectivism; the other side of the coin is that, since almost philosophical views in this domain recognize this multiplicity, there is no easy route from the multiplicity intuition to any substantive philosophical position. In the literature on perspectivism, there are two arguments for the position that build on the multiplicity intuition: the argument from incompatible models and what I am call the argument from failure of fit. The argument from incompatible models has received the most detailed discussion in the recent literature, so I will begin here. (This argument is important in both Giere and Teller; see also Chakravartty 2010; Morrison 2011, 2015; Rueger 2016 for discussion.)

The rest of this chapter proceeds as follows. I begin by outlining the argument from incompatible models, and then turn to the different forms of perspectivism that have been proposed in response to this argument. Can the perspectivist offer something beyond traditional realism and antirealism, or do these perspectivist views ultimately collapse into one of these more traditional positions? I will argue that only a relativist perspectivism constitutes a distinctive position. I will then evaluate whether the argument from incompatible models provides a good justification for this position.

4.2. The argument from incompatible models

Cretu (2018) states the argument as follows:

(P1) Diverse modelling practices are ubiquitous in contemporary scientific practice.

(P2) Diverse modelling practices often give rise to multiple, different models of the same target system.

(P3) Different models constitute partial perspectives regarding the same target system.

(P4) Knowledge obtained through partial perspectives regarding the same target system is inherently perspectival.

(C) Therefore, modelling practices yield partial and perspectival knowledge about the world.

This is a synchronic argument, appealing to facts about present scientific practice. Essentially the same argument can be run using theories and models from the history of science, giving a diachronic incompatible models argument. The appeal to the history of science is a classic objection to scientific realism; different forms of argument are discussed in Wray (2015). Here is Cretu's perspectivist formulation of it:

(P5) Once we reflect on the succession of past scientific theories we are led to the conclusion that although many past theories proved to be false, some at least, got something right about the world.

(P6) Since many such theories have been abandoned, yet some claims were retained, it cannot be the case that science gives us true and complete images of the world.

(P7) Instead, what history of science teaches us is that there are multiple different perspectives regarding the same phenomena.

(C1) Thus, history of science yields perspectival knowledge of the world.

We can yield a general incompatible models argument by replacing (P1) of the synchronic argument with:

(P1*) Diverse modelling practices are ubiquitous in *successful* scientific practice.

Where "successful" now covers not just successful contemporary theories, but also those successful past theories that have either been displaced by better theories, or that are still retained but taken to be false, such as Newtonian mechanics. This is an addition that we need to make in any case, because the traditional debate between

realism and antirealism has always concerned specifically successful scientific practices. Of course, what is meant by “success” here is open to interpretation, and different realists would elaborate on this in different ways. The important point for now is that whatever criterion we use to judge contemporary modelling practices successful, it is very likely that some past models will come out as successful as well.

Past theories that we now know to be false did get *something* right about the world, and in this respect, they are analogous to partial, idealized models. Thus, whether we take the proliferation of incompatible models in contemporary science, or the “graveyard of failed theories” (Laudan 1981) from the history of science that nevertheless managed to latch onto regularities in the world, we are faced with a variety of incompatible representations that seem to get something right. (The connection between perspectivism and Laudan’s pessimistic induction is discussed in detail in Fellowes 2020.)

One point to note about the incompatibility argument is that incompatibility might be found at different “levels”. Scientists might construct incompatible models of a phenomenon from the same theoretical perspective. Or they might adopt alternative theoretical perspectives – alternative theoretical principles, which can then be used to construct different models. Giere talks about the “Newtonian perspective”, the “Darwinian perspective”, and so on. Since on Giere’s view, such theoretical principles do not in themselves make any claims about the world, he would not describe this situation in terms of incompatibility. Of course, when attempting to describe the underlying structure of the world, we would no longer use the Newtonian perspective. A standard realist would probably take the Newtonian perspective as straightforwardly incompatible with our own. In any case, the point for now is just that there is a distinction here. We may construct incompatible models from the same broader theoretical perspective. Or we may have two different theoretical perspectives, leading to the construction of alternative models.

The obvious question concerning the argument from incompatible models is: why perspectivism? This same argument has already been used in defence of various antirealist views; and for the historical argument specifically, there is an enormous

literature on the challenge that successful theories of the past pose to realist interpretations of science. When faced with incompatible, successful models, we might simply conclude that these models fail to represent reality. The key motivation for perspectivism is preserving a form of realism in the face of incompatible models. Realists hold that our best theories aim to give us true descriptions about what the world is like. Yet scientists frequently appeal to incompatible models of the same system (Morrison 2011). We cannot take these models as simply being true, or we would be committed to contradictory accounts of the way the world is. What are our options? Prima facie, it seems we must either deny that the models can be interpreted realistically; or we hold that the world is contradictory, literally containing incompatible states of affairs; or we embrace Goodman's (1978) claim that there are literally many actual worlds. Neither of the latter two options are attractive to most realists. Perspectivism seemingly offers an alternative.

Perspectivism, it is claimed, holds the potential for vindicating a form of limited realism. Where exactly is the problem for the realist? The issue, I take it, is that realists want to draw a connection between success and truth. Many realists, following the "no miracle" argument, will take truth as the explanation for success, so that realism is presented as the conclusion of an inference to the best explanation about the remarkable success of science. Those realists who resist IBE, or at least attempt to, will still see success and truth as importantly connected. We find this in Hacking's entity realism (Hacking 1983) and Hofer's recent "tautological realism" (Hofer 2019). In these cases, there is arguably a move away from IBE, but still, it is the success of science, broadly speaking, that indicates its truth. Hacking emphasizes our success at constructing reliable instruments; Hofer emphasizes success at unifying various aspects of theories from a host of different domains. The general challenge, then, is that reflection on scientific practice undermines the motivation and justification for realism. The realist holds that success indicates truth. Kitcher calls this the "success-to-truth rule":

S plays a crucial role in a systematic practice of fine-grained prediction and intervention.

So, S is approximately true.

But then the realist is presented with a host of models that exhibit what she takes to be the relevant kind of success, but which, it is proposed, cannot all be true. The success of such models must be explained by something else: perhaps we can explain the success of Newtonian models by appealing to the truth of *other* theories and models to which the Newtonian models relate in the right kind of way. Assuming that general relativity is true, we can show how to recover the Newtonian laws as approximations in limited circumstances. But now the problem is: why shouldn't the same be true for those theories and models that have not yet been superseded? Perhaps there is some other, as-yet-unconceived theory T, such that if we accepted T, we would take general relativity to be false, but we could use T to show how general relativity delivers approximate truths in specific circumstances.

An interesting point made by Elgin (2019) is that there is something rather strange about the incompatible models argument as a challenge to realism. The trouble arises not because our theoretical options are limited – it is not that we have failed to hit upon the best theory or model – but rather that we have an “embarrassment of riches.” There are too many good models, too many good theories: too many theories that are empirically successful and that strike a good balance on other theoretical virtues. In this situation, we are unable to select one single model as correct. Even in ideal circumstances – however we specify what exactly counts as “ideal circumstances” – there is no reason to suppose that our theoretical virtues will deliver a unique solution. There might always be some other model, perhaps as yet unconceived, that does just as good a job or better (cf. Stanford 2006).

Facing two equally adequate alternative models, traditional realists take it that only one can be true. Since we have no grounds for deciding between them, we are driven into skepticism. Either M1 or M2 is correct, but not both; we cannot know which; so we cannot know the way the world is. What's odd about this situation, as Elgin notes, is that if we had only T1 on the table, we would be willing to accept. It is the abundance of good theories that leads to skepticism.

How is perspectivism supposed to help? The general idea is that apparently incompatible claims must be understood as perspectival; just as “Frank is to the left of Vincent” and “Frank is to the right of Vincent” can both be true, provided we

specify the perspective from which we are observing Frank and Vincent, so it is with scientific claims. If knowledge is relative to perspective, then there is not really any incompatibility between apparently inconsistent claims. We need not endorse two inconsistent propositions; it is rather that the models provide different perspectives on the phenomena. As Rueger (2016: 402) puts it:

Drawing on the analogy with different spatial perspectives onto the same object, the perspectival realist suggests that the seemingly incompatible representations given by our models are in fact just representations of the same (independently existing) target system from different perspectives or points of view.

Apparent incompatibility is removed by treating the alternative models as perspectival representations. We do not simply represent the system; we represent the system from a perspective. So, it is no surprise that the representations are different. We can take it that both models are true, just true-from-a-perspective. The details of how this is supposed to work depend on the form of perspectivism in question. I think it is useful to distinguish the more moderate forms of perspectivism that attempt to avoid alethic relativism (i.e. truth relativism, the view that truth is relative to perspectives) from the more radical forms that embrace it. I will discuss both approaches in this chapter.

Before outlining these forms of perspectivism, it is worth keeping in mind that critics of perspectivism can raise two questions, which it is important not to conflate. First, is there a distinctive perspectivist position? Is it possible to adopt a stance on incompatible models, other than straightforward realism or instrumentalism? Corresponding to this, the critic may argue against perspectivism by trying to show that it is not a stable position. Under analysis, the positions outlined in the next part will, it is suggested, collapse into either realism or instrumentalism. The second question is, assuming that perspectivism is a distinctive position, should we adopt this position? Corresponding to this, the critic of perspectivism may argue that while it is a possible position, there is simply no good reason to adopt it.

Suppose we have two incompatible models, M1 and M2, both of which are extremely successful. There is no doubt that we can adopt one of two positions in response to these models. We might take the traditional realist line and try to identify one of these models as accurate or try to select specific claims from both models that can be taken as true without any incompatibility. Or we might be instrumentalist. Now, is there a viable third option here? Part of our task is to spell out what this third option could be. But then there is the additional question of what the justification is for adopting this third option. Both perspectivists and their critics have, I think, sometimes conflated these issues: the appeal to incompatible models is sometimes presented as demonstrating the coherence of perspectivism; and sometimes presented as an argument for perspectivism.

4.3. Moderate perspectivism

I begin with moderate perspectivism, which has been elaborated by Giere, at least under one interpretation of his work, and by Massimi. The general idea of the argument from incompatible models is to undermine traditional accounts of truth. This can be developed in two ways. First, the moderate perspectivist might downplay truth and understand the model-world relation in terms of similarity (or something else such as isomorphism, etc.); the perspectivist then holds that other forms of realism and antirealism go wrong in viewing scientific knowledge in terms of truth. Second, the moderate perspectivist can accept truth but develop some notion of “perspectival truth”. Giere takes the first option; Massimi, the second.

4.3.1. Similarity

The similarity approach is developed in Giere (2006). Giere argues that models represent in virtue of similarity relations: scientists use models to represent systems in the world by exploiting similarities between the models and those systems (2006: 63). No model ever perfectly fits its target, so scientists must specify the required respects and degrees of similarity. The required degrees of similarity will depend on our goals, our computational power, and the accuracy of our instruments, so in model-building we can never achieve a view-from-nowhere. These similarity comparisons are expressed in what Giere calls “theoretical hypotheses”, so a

theoretical hypothesis will tell us something along the lines of: “model M is similar to system S in respect R to degree D.” What is important for the perspectivist is that this is supposed to remove the central role of truth in scientific representation.

Models are nonlinguistic, so cannot be true or false. As Giere puts it:

the empirical representational relationship is not directly between statements and the world, but between models and the world. Here the operative notion is not truth, but similarity, or ‘fit’, between a model and the world. (Giere 1999: 73)

Two key aspects of fitting models to the world are interpretation and identification (Giere 1988: 72; Giere 2006: 62). Interpretation involves connecting the mathematical symbols in the model with general concepts that are supplied by the scientific perspective, such as how we may use the Newtonian perspective to interpret certain mathematical models in terms of *force* and *position*. The perspective defines the concepts and relates them to other concepts, such as *mass*, *velocity*, *acceleration*, etc. Identification involves connecting the mathematical symbols, or the non-mathematical elements of the model, with specific things in the world, such *the position of the Moon*. Here is another example. Take the simple model of exponential population growth:

$$dN/dt = rN$$

The interpretation is as follows: N is the population size; dN/dt is the change in N over time; r is the reproduction rate minus the death rate. Then to apply this model to the world, we must identify elements of the model with things in the world, so that we may identify N with, for example, a real population of deer released on an island, and then we would identify r with the reproduction rate minus the death rate of this particular population. In this case, the relevant similarity is the growth rate of the population. The real population will soon depart from the predictions of the model, since the model population increases exponentially forever. So, we can develop more sophisticated models, e.g.:

$$dN/dt = rN((K-N)/K)$$

where K is the carrying capacity of the environment. In this model, the population plateaus. If N is identified with a real population of deer, then K is identified with certain properties of the environment in which the deer live. Thus, interpretation and identification allow us to generate representational models that can be linked with real-world systems. The models represent specific aspects of the systems, and only more-or-less well, never perfectly.

Similarity provides the grounds for a perspectival solution to the incompatible models problem because there seems to be no contradiction in endorsing all the following claims:

- M1 is similar to S in respect R1.
- M2 is similar to S in respect R2.
- M1 and M2 are incompatible.

Or even:

- M3 is similar to S in respect R1 to degree D1.
- M4 is similar to S in respect R1 to degree D2.
- M3 and M4 are incompatible.

“Incompatibility” here would be understood in the sense that the models provide different pictures of the world. Take a model of a star which attributes to it constant density, and a model of a star which takes its density to vary linearly with its radius. These are incompatible in the sense that if both were taken to fit a given star perfectly, then we would have to attribute contradictory properties to the star: its density would be both constant and not constant throughout. But perfect fit is never available with models of physical systems. Our two stellar models may represent the structure of the convective and radiative zones in the star, and they may represent them as having different properties; but since the model-users only aim for similarity between the model and the system, both can be taken to be adequate. So, this appears to offer a straightforward solution to the incompatible models problem. As Teller (2016) puts it:

Chemists know that, when it comes to chemical phenomena, the world is very like one populated by atoms and molecules, Quantum field theorists know that, when it comes to quantum field theoretic phenomena such as superposition of different numbers of quanta, the world is very like one filled with a quantum field.

Presumably, the same idea can be applied in the case of diachronic incompatible models. Giere's discussion of the diachronic incompatible models argument is brief, at less than a page. He states the argument in form of the pessimistic induction: most past theories were shown to be false; by induction, our present successful theories will likely be shown to be false too. Giere then claims that the pessimistic induction ought to be rejected because it "assumes the older standard view that scientific theories are sets of straightforward empirical statements that must be either true or false" (2006: 95). As we have seen, on Giere's view, models are the primary representational tools, and models cannot be true or false, because they are nonlinguistic. Only theoretical hypotheses, which specify the fit between models and the world, can be true or false; and this fit is never claimed to be perfect, in any case.

To elaborate a little further on this, even failed theories may allow us to construct models that represent aspects of real systems with a reasonable degree of accuracy: for example, using caloric theory, scientists were able to construct models of phenomena such as the state changes of matter, and we would still judge these models to be similar to the target systems in significant respects, if only in terms of their representation of observable properties. To this day, we still construct Newtonian models of various systems, even though the theory is no longer taken to correctly fit the underlying structure of spacetime. There are good practical reasons for adopting the Newtonian perspective, and its models often exhibit a high degree of accuracy in the relevant respects. However, we no longer think of the world itself in terms of the interpretations that the Newtonian perspective provides.

There are several difficulties with this approach, and ultimately, I do not think it provides a genuine alternative to more standard accounts of realism and instrumentalism. So, keeping in mind our two questions, we need not ask whether

there is good justification for adopting this perspectivist alternative, because in fact, there simply is no coherent alternative position here, in the sense that there is nothing here that adds anything new to existing accounts. There are two reasons why taking similarity as the primary representational relation cannot support a *via media*.

First, note that realists already frame their view in terms of approximate truth (Boyd 1983) or truthlikeness (Niiniluoto 1999). Now, approximate truth is not identical to similarity – though arguably, approximate truth is one kind of similarity. Perhaps it would make little sense to speak of an approximately true theory as being “similar” to the world, but surely, an approximately true theory would in important respects be similar to the exactly true theory. As a result, there is little surprise in realizing that approximate truth does all the work that similarity is supposed to do in terms of accounting for incompatible models. Take “M1*” to refer to the set of propositions associated with a model. The standard realist can endorse all these propositions:

- M1* is approximately true of S in respect R1.
- M2* is approximately true of S in respect R2.
- M1* and M2* are incompatible.

Earlier I spoke of models, now I speak of “propositions associated with models”. Is this move legitimate? Approximate truth cannot do all the work of similarity, since it cannot attach directly to models, at least assuming that Giere is correct that models are nonlinguistic. Suppose we have a model of the Sun, which depicts it with a core in which thermonuclear fusion takes place, surrounded by a radiative zone, surrounded by a convective zone. In the model, the core is about 0.2 of the solar radius. Any similarity that holds between the model and the Sun can instead be stated as a proposition which is approximately true. If the internal structure of the star in the stellar model is similar to the internal structure of the Sun, then “the Sun’s core is 0.2 of the solar radius” is approximately true.

One of Giere’s motivations seems to be something like this: If we are talking in terms of truth and propositions, then we will have to say either that there is one true view of the world, or that all views are just false. Similarity is supposed to give us a way of

accepting a range of different views; different ways that the world appears from different perspectives. The trouble is that approximate truth can play the same role.

Indeed, note that defenders of the pessimistic induction, and other forms of incompatible models arguments against realism, will frame their objections so as to accommodate the realist's appeal to approximate truth, and these arguments, if successful, will also undermine similarity. Phlogiston theory was successful in many respects in its time, but "combustion releases phlogiston" is not even approximately true. The miasma theory of disease supported a variety of successful novel predictions (Tulodziecki 2016), but it is thoroughly mistaken in its description of the transmission mechanism of disease. These theories fail, whether we look at them in terms of approximate truth or similarity. They completely misrepresent how the world works. That is the antirealist argument, at least. Shifting to similarity changes nothing: any examples used to undermine approximate truth will undermine similarity just as well. So, the point that models never exhibit perfect fit simply misses the mark as a response to the incompatible models argument. There is no distinction from standard realism here.

Along the same lines, problems of incompatibility re-arise in terms of similarity. What happens when there are incompatible similarity claims? Suppose we have:

- M1 is similar to S in respect R to degree D.
- M2 is similar to S in respect R to degree D.
- M1 and M2 are incompatible.

This is the same problem as the original incompatible models argument. This is a problem that occurs with incompatible models regardless of how we characterize the representational relation. We can model water as a continuous fluid or as a collection of discrete particles, and the resulting models might be similar in certain respects, but there is no obvious way of taking both the continuum and particle pictures as similar to the real underlying structure of water. Of course, the perspectivist could suggest that we need to think of *similarity* in an alternative way, a way that is different from the standard realist understanding of similarity relations, and that captures how similarities are a product of the interaction between scientists and the

world... and so on. But this undermines the point of introducing similarity in the first place, which was to resolve these problems as they arise with the notion of truth.

The only way out of this is to adopt an account of the representational relation in which the incompatibility cannot be stated. For example: to say that X and Y are incompatible is to say that X and Y are contradictory; but contradiction is a property of propositions; and models are not propositional. So we never really have incompatible scientific models. But unless we want to deny that models represent the world at all, we will find ourselves needing to re-describe incompatibility in different terms. Again, the issue is that there is no obvious way of viewing water as similar to both a continuous fluid and a collection of particles, at least if we are talking about the same respects and degrees. Attempting to imagine this is akin to attempting to imagine a round square.

There are, of course, other motivations for the similarity view. Giere also suggests that a further problem with approximate truth is that nobody has ever given a good account of what it consists in: in this respect, we have a much better grasp of similarity than we do approximate truth. I know what it is for a model to be similar to a target system; it is much more difficult to say what it is for a set of propositions to be approximately true of some target system:

First, there is the spotty history of attempts to give a formal account of the logical properties of approximate truth (Oddie 2014); there is no similar history of failure of formal accounts of similarity. But this is surely just a result of the fact that no attempts had been made to give such formal accounts. Indeed, if anything, approximate truth seems to be on better ground here: we do in fact have a good idea of where to start, we have well-developed formal theories, and we know what their problems and prospects are.

Second, it might be argued that approximate truth involves a host of other philosophically problematic notions: for example, Laudan (1981) supposes that approximate truth of a theory requires that the theoretical terms of the theory refer, but reference is also the subject of a great deal of debate. What exactly is the reference relation? Under what circumstances do theoretical terms refer? A benefit

of the move to similarity is that this problem evaporates: reference of theoretical terms is obviously not required for similarity between a model and target. However, my suspicion is that counterparts of all the philosophically controversial notions associated with approximate truth will be found for similarity. To use models as representational tools, we must *identify* elements of the model with particular things in the world. Identification is the counterpart of reference. Surely, for a model to be similar to a target system, there must be an identification of elements of the model with elements of the system. For one thing, I must be able to identify the model with the system, otherwise I can't say that it's a model of that system! My conclusion is that approximate truth does all the work that similarity can do, and similarity does not offer any clear solution to the problems facing approximate truth.

The second problem with Giere's appeal to similarity is that similarity is parasitic on truth. This point has been made by Chakravartty (2001). In Giere's case, we see this with the introduction of theoretical hypotheses, propositions which state the respects and degrees of similarity between models and the world. As such, they are either true or false. So, judgements of truth and falsehood are unavoidable at some point in the analysis. But there is a deeper problem here. The similarity judgements expressed in theoretical hypotheses are either totally uninformative, or if they are framed in an informative way, they will be eliminable in favour of more precise propositions which do not involve similarity claims.

What is it that we are doing when we make a similarity judgement? We are relating properties of the model to properties of the worldly system. But simply to say that X is similar to Y, or even that X is similar to Y in a particular respect and to a particular degree, tells us nothing. Take the wallet on my table. Suppose I tell you that the wallet is similar to an elephant. Perhaps this is an appropriate claim. But clearly, in the absence of a specification of the properties that the wallet shares with elephant, we learn nothing from it. This remains the case even if I tell you *in what respect* the wallet is similar to an elephant an elephant. Consider: the wallet is similar to an elephant with respect to its material constitution. What does this mean? Here are some things that I might be saying:

- The wallet is made of parts taken from an elephant, so the wallet is similar to an elephant in the sense that the wallet is made of elephant cells, that certain patterns on elephant skin can be seen on the wallet, etc.
- The wallet and the elephant are both made of chemical elements.
- The material of the wallet is in the same state of matter as the material of the elephant.
- The material of the wallet is approximately as effective at stopping bullets as is elephant skin. That is, you would need approximately the same thickness of wallet material, and of elephant skin, to stop a bullet.
- The same philosophical problem of material constitution arises with respect to both the wallet and the elephant: the wallet (elephant) and wallet-material (elephant-material) are both coextensive, yet they seem to have different persistence conditions, because if I crushed the wallet (elephant), I would destroy the wallet (elephant), but I would not destroy the wallet-material (elephant-material).

This can, of course, be continued almost indefinitely. The point I am making here is not Goodman's (1972) concern that anything is similar to anything else in an infinite variety of respects. Perhaps there is a sensible distinction to be drawn between intrinsic or "natural" properties, such as an object's shape, and relational properties, such as the property of being located more than 1 million miles from the Sun. Certainly, I do not object to making similarity judgements. The point is simply that in judging any two objects as similar in some respect, you are comparing specific properties of those objects. If you fail to specify the properties, the similarity judgement fails to communicate anything about either object and will be useless in the context of the sciences; but if you do specify the properties, then you have made a statement which is either true or false. Of course, this is one reason why Giere introduces theoretical hypotheses in the first place.

I think this is made especially clear when we consider the various models that are useful precisely in virtue of dissimilarities: the Hardy-Weinberg model tells us what happens to allele frequencies in the absence of evolutionary forces, so by studying how real populations *deviate* from the model, we can conclude that evolutionary forces are at work in these populations. Yet it is also true, and in certain contexts would be entirely appropriate, to say that the population in Hardy-Weinberg

equilibrium is similar to a real population that is not in Hardy-Weinberg equilibrium. It is simply that the model population and the real population share some properties but not others.

What all of this means is that Giere's moderate view is open to a straightforward application of Escape from Perspective. Whenever a similarity judgement is made, the question is: similarity between what? We have our theoretical hypothesis:

(P) Model M is similar to S in respect R to degree D.

Perhaps we should think of M as providing a perspectival representation. But what of (P)? It's either true or false in a straightforward sense. Worse, it can only be evaluated for truth or falsehood once we specify the properties of the model and the system in question, but having done this, similarity no longer plays any significant role. The kind of perspectivism involved in Giere's view here amounts to nothing more than the view that the representations afforded by models inevitably present things from the point of view of the model. This is hardly troubling for a standard realist, if we can stand outside the model and seeing how it relates to reality. But this is exactly what we will do, in describing the properties of models and properties of target systems. In this way, similarity judgements appear to require the application of Escape from Perspective. To say that X and Y are alike in a particular respect involves taking a stance outside of both X and Y and comparing them – a kind of “god's eye view”. We say that the map is similar to the territory, because we can observe both the map and the territory independently.

All we are left with in the similarity account as an alternative to standard realism is Giere's claim that similarity avoids problems associated with the correspondence view of truth, and thereby provides a better basis for a naturalistic realism. I cannot see how similarity between an abstract model and a physical system is any less mysterious than correspondence between a proposition and the world; after all, models and worldly systems are completely different kinds of things, so there are difficulties even making sense of the claim that, say, a stellar model has similar density to the Sun (cf. Thomson-Jones 2010). But even if such difficulties could be solved, there is nothing here that would constitute an alternative to standard realism.

Indeed, if anything, it seems to me that the addition of similarity as a tool of representation makes for a stronger realist position than talk of truth alone. Part of the reason for this is that whereas there are now a variety of accounts of truth, many of which are compatible with realism (cf. Musgrave 1989 on how realists may accept deflationary accounts of truth), it is hard to see how to understand similarity if not in terms of a kind of correspondence: certain features of *this* thing correspond to certain features of *that* thing.

4.3.2. *Truth*

I noted that the incompatible models argument prompts one of two responses from the perspectivist: either downplay truth and understand the model-world relation in terms of something else such as similarity, or develop a notion of “perspectival truth”. I have argued that Giere’s similarity approach does not support a distinctive perspectivist position. Perhaps there is some other representational relationship that would do the job here. However, as far as I am aware, nothing has yet been proposed in the literature by any perspectivist; and I think it is clear enough how the arguments given above could be applied to other views of representation such as isomorphism, were a perspectivist to appeal to these instead. So, I will not pursue this option any further.

An alternative approach to perspectivism, which is moderate in that it explicitly attempts to resist relativism, has been developed in the most detail by Massimi. As noted in the introduction, Massimi takes “perspectivism” to be a thesis not about what we have scientific knowledge of, but rather about *how* we form scientific knowledge (Massimi 2012). The truth-makers for scientific claims are mind-independent, perspective-independent facts about the world; justification of scientific claims is perspectival. The question raised by incompatible models is how we can justifiably take any claims to be true. Massimi’s goal is to save a form of limited realism in the face of epistemic pluralism. Scientific research has generated a variety of competing theories and models for explaining phenomena, both across history and in our own contemporary scientific practice, but we still want to claim that these theories and models yield knowledge of the way the world really is, independently of our perspective.

Massimi (2016a) summarizes her account of perspectival truth as follows:

Knowledge claims in science are perspective-dependent when their truth-conditions (understood as rules for determining truth-values based on features of the context of use) depend on the scientific perspective in which such claims are made. Yet such knowledge claims must also be assessable from the point of view of other (subsequent or rival) scientific perspectives.

The truth-conditions for scientific claims are standards of performance adequacy that are fixed within a given scientific perspective. Scientific perspectives lay down standards of performance adequacy for the claims made within those perspectives. For example, standards that are shared by almost all scientific perspectives are that their models be consistent with experimental data and that their models allow us to make projectable generalizations. Other standards are more specific to the problems of the time. In 19th century physics, models of the ether were developed to account for the phenomenon of electric displacement. So one of the standards of performance-adequacy of this perspective was: furnish a good explanation of electric displacement. But models satisfying all these standards of performance-adequacy could not be developed (Massimi 2016a: 15).

Massimi then argues that not only do scientific perspectives specify standards of performance-adequacy for their own claims, they also function as contexts of assessment for evaluating the claims of other perspectives. In particular, I can ask, from the point of view of scientific perspective SP1, whether claim P of scientific perspective SP2 satisfies the standards of performance-adequacy of SP2. I can evaluate, using the contemporary electromagnetic perspective, the degree to which 19th century ether models satisfied the standard of consistency with experimental data, a standard which was laid down by the earlier 19th century perspective. With this in mind, Massimi's account of perspectival truth is this: Some scientific claim P is true-within-perspective-SP, just in case (1) P is true, in the sense that P corresponds to a mind-independent fact, and (2) P satisfies the standards of performance adequacy laid down by perspective SP, *when assessed from other perspectives SP1, SP2, SP3, etc.*

Through condition (1), scientific claims are made true or false in virtue of mind-independent facts about the world. P is either true or it isn't, regardless of what anybody thinks about it. This still leaves us with the question of how we can ever be justified in taking any particular claim to be true. While most realists will simply apply a success-to-truth rule, this runs into the problem of the proliferation of different perspectives on the world, which represent the world in different ways and all have different standards of success. So, whose success matters, and how do we evaluate success? If we take it that *our* standards of success are what matter, then we will judge that the standards of previous perspectives are defective, and do not indicate truth. But if we grant that science may improve in the future, then from the point of view of future scientists, our standards will similarly count as defective for them, and from their epistemically superior position, our standards will not indicate truth. On the other hand, suppose we permit the success-to-truth rule to be applied, per whatever the standards of success of a given perspective are. Well, now we face the standard incompatible models argument. Scientists of the past would apply this rule to conclude that their descriptions of phlogiston, caloric, and the luminiferous ether, were true.

This is where condition (2) plays an important role. Our justification for thinking that P is true or false is tied to scientific perspectives. While a scientific perspective will specify standards of performance-adequacy, this is evaluated from other perspectives. The success-to-truth rule is applied when we judge that P meets the standards of performance-adequacy of its perspective, from the point of view of other perspectives. In this way, we can pick out the successful aspects of past theories, without assuming that there are certain "eternal" standards of success, without assuming a kind of "view from nowhere". We simply use our present perspectives to assess whether the claims of past perspectives were successful by the standards of those past perspectives. Furthermore, by requiring that any claims true-within-a-perspective must be positively evaluated from other perspectives, we avoid the problem of perspectives vindicating themselves, and of having to applying the success-to-truth rule by whatever the standards of success of a given perspective happen to be.

As Massimi makes explicit, this is a kind of selective realist strategy (2016b). The general idea of selective realism is that, faced with the problem of historical discontinuity in the science, we must distinguish the “working posits” from the “idle wheels” of our best theories, where the “working posits” of a theory are those elements of the theory that were used in producing its successes. The selective realist’s claim is that whereas the idle wheels are thrown out over the development of science, giving us the historical discontinuity, the working posits are retained. The same idea can be applied to the problem of incompatible models more generally. When faced with incompatible models, both of which are empirically successful, we must identify which aspects of the models are involved in the derivation of successful predictions; and hopefully, there will be no incompatibility between these elements. But how do we distinguish the working posits from the idle wheels? How do we tell whether past theories were successful in relevant respects? Perspectives define standards of performance-adequacy, and they provide contexts of assessment for judging whether claims of other perspectives meet their standards of performance-adequacy. As Massimi puts it:

Scientific perspectives ... provide *contexts of assessments* for scientific claims. Galileo could assess the Aristotelian claims about free fall and find them lacking in satisfying what, from Galileo’s own perspective, were the standards of performance adequacy *appropriate to the Aristotelian epistemic context*. (Massimi 2016b: 105)

Does Massimi’s position provide a genuine alternative to standard realism? Here is a concern. On Massimi’s account, in order for some claim P to be successful, or for P to count as perspectively-true, it must be the case that P is true, period. P is true-in-a-perspective when P is true and is judged to meet certain standards of performance-adequacy. The trouble is that in identifying what counts as true, we are straightforwardly using our own scientific perspective, and hence run right back into the problem that we must be taking our standards of success as the “eternal” standards. How else could we determine what the truth is? We can only use what we take to be our best theories. But we only take those theories to be true because those theories meet our standards of success (*if we take them to be true, of course*).

The realist applies a success-to-truth inference. Success indicates truth. Massimi's concern is that different perspectives have different standards for success. Now, most realists assume that our standards of success are what indicate truth, but the problem with this view – the problem with supposing that it is, as Massimi puts it, “success from here now” that indicates truth – is that both scientists of the past, and of the future, can apply the same reasoning. Past scientists conclude that various things are true which we know to be false, and future scientists conclude that *our* beliefs are false and our standards are defective. So, Massimi proposes that perspectives function as contexts of assessment for the claims of other perspectives. In this way, we can see how various claims are reasonably retained over time, and were reasonably taken to be true by scientists of the past. Suppose that our perspective is SP1, and we are evaluating some claim P from SP2. P only counts as true-within-SP2 if P is true. Now, why do we believe that P is true? Well, because P is a part of our own best theories, and those theories are successful according to our own standards for success. Ultimately then, there is no way in which this alters the standard success-to-truth inference.

Let me state the same problem, but from a slightly different vantage point. Part of Massimi's goal here is to save a form of realism in the face of scientific theory change. There are various propositions that modern scientists endorse. Take one of these propositions, P. How can we be justified in taking P to be true? Massimi's perspectival realism will use contemporary perspectives to assess the performance of propositions in past perspectives, to “identify scientific claims that – by being *justifiably retained* in the shift from the original [perspective] to another [perspective] – we have reasons for thinking of as true” (2016b: 108). So, it seems that what justifies us in believing that P is true has something to do with its performance in past perspectives. But this is not the case, and it seems that Massimi explicitly accepts that this is not the case. Later in the same article, Massimi considers the example of scientists who get things right, but who fail to satisfy the prevailing standards of performance adequacy (2016b: 112). She cites the astronomer Vesto Slipher, who discovered the expansion of the universe a decade before Hubble. Slipher was working at a time when the universe was assumed to be static, before models of an expanding universe had been properly developed. It seems plausible enough that Slipher's claim was unsuccessful. However, we would certainly not say that Slipher's

claim was false. Our reason for thinking that he was right, that it is true that the universe is expanding, has nothing to do with the performance of this claim in past perspectives. We use our own standards of success to determine what is true, and then we use that to determine what was true for scientists working in the past. Condition (2) is not actually doing any work here.

Could we drop condition (1)? In that case, we would take a claim to be true within a perspective, provided it satisfies the standards of performance adequacy laid down by that perspective, when assessed from other perspectives. This does have the benefit of being a genuine alternative to standard realism, and it also allows for a straightforward response to the problem of incompatible models. “Water is a continuous fluid” may satisfy the standards of performance-adequacy of perspective SP, while “water is composed of discrete molecules” satisfies the standards of performance-adequacy of perspective SP*; in different contexts, it may be useful to make different assumptions about the underlying nature of water. The obvious problem here, for the attempt to stake out a moderate perspectivist position, is that dropping condition (1) gives us a straightforward relativism about truth. P is true within SP provided we assess it as meeting the standards of performance-adequacy within SP. In that case, even from our own perspective, any proposition P could meet the standards of performance-adequacy of some scientific perspective, and thus count as true-within-that-perspective, if that perspective makes the standards low enough. This is not a merely theoretical concern. After all, we can see, from our perspective today, that postulating aether, phlogiston, and caloric, and so on, satisfied certain standards of earlier perspectives.

In fact, dropping condition (1) would not solve the fundamental problem, because there is a further difficulty in Massimi’s account. How do we tell what the standards of performance-adequacy of a given scientific perspective are? Take the standard of accuracy to the data, a central standard for pretty much any perspective that we would recognize as scientific. To determine whether some claim meets this standard, we must first specify: accuracy to what degree? As Wray (2018: 165) points out, the standard of accuracy changes over time. In the medieval period, astronomical theories were expected to locate planets with an accuracy of within about 15 degrees of arc. By the time of Copernicus, this had been raised to an accuracy of

about 5 degrees; and then Tycho Brahe's innovations in astronomical observation raised the standard of accuracy significantly further. Consider the orbit of Venus within the Ptolemaic system. Does the Ptolemaic model of the orbit of Venus meet the Ptolemaic perspective's standard of accuracy to the data? Well, the question is how exactly this standard is to be understood. Here are two options:

- The Ptolemaic perspective imposes a standard of adequacy to all observations of the sky, to within approximately 5 degrees of arc.
- The Ptolemaic perspective imposes a standard of adequacy to all possible observations. (That is, it is adequate no matter what instruments we use or how precise our data is.)

Certainly, the Ptolemaic model of the orbit of Venus was considered successful at the time. Today, using our contemporary astronomical perspectives, we judge that it did indeed meet the standard of accuracy to within 5 degrees. But with the far greater precision allowed by our observations, we would now judge it to fail to meet the standard of accuracy to all observations in general. So, for Massimi's approach to be plausible, standards of performance-adequacy need to be specified within perspectives in a fine-grained way. Unfortunately, it seems to me that there is no obvious reason for favouring either interpretation. The standards of performance-adequacy of a perspective are, I think, often indeterminate. Part of the problem here is that standards of performance-adequacy are laid down not by perspectives themselves but rather by scientists. After all, we still use the Ptolemaic perspective to this day in some contexts, but the standards of performance-adequacy would be very different, since we do not put it to the same uses as earlier astronomers – for one thing, we no longer treat it as a providing true descriptions or explanations of celestial phenomena.

Even putting this point aside, and considering scientists at a specific place and time, specification of the standards of performance-adequacy is extremely difficult. There is a sense in which Ptolemaic astronomers worked with a standard of accuracy to 5 degrees of arc. This was the best they could do, given the limitations of their instruments. However, there is also a sense in which any perspective that is recognizable as a *scientific* perspective must in some sense require adequacy to all

possible observations. One thing we that any philosophy of science must account for is the obvious fact that scientists are constantly trying to seek out new observations and to improve their observational accuracy. If we take it that the Ptolemaic perspective required accuracy to within 5 degrees, then provided that its models met this standard, what would be the point in astronomers attempting to make new observations, to within a greater degree of accuracy? They already had models that were empirically adequate per the Ptolemaic standard. Any further observations would not be capable either of lending further support, or of raising any challenge, to these models, when interpreted within this perspective. But the other side of the coin here is that if we require that, to be true-within-a-perspective, a scientific claim must meet the standard of adequacy to all possible observations, no matter how precise, then very few (perhaps zero) currently accepted scientific claims will turn out to be true-with-a-perspective, and we will be left only with claims that we take to satisfy condition (1).

The standards of performance-adequacy that we attribute to past perspectives is strongly dependent on the perspective we are taking on those past perspectives. We might ask a historical question, about how scientists developed models that were adequate the observations they had available to them, or about how they built instruments to reveal new phenomena, and the impact this had on the evaluation of the models. In this case, we will think of the Ptolemaic perspective as imposing standards of performance-adequacy of within 5 degrees of arc. Alternatively, we might ask about the rationality of the scientific process more generally, in which case we will think of the Ptolemaic perspective as imposing much more stringent demands.

There is a final point to be made about Massimi's position. As noted, Massimi's perspectivism is epistemological. The point of perspectivism, for Massimi, is to provide a more plausible account of *how* we develop and justify scientific knowledge than had been achieved by earlier realist approaches. But Massimi makes no room for perspective-dependent facts. For some proposition P to count as true-within-a-perspective, it must be true in the sense that it corresponds to a mind-independent, perspective-independent fact. Justification is perspectival; the facts are not. But epistemological perspectivism is perhaps tangential to the type of perspectivism that

is proposed in Giere, and that is the focus of this thesis. It is worth comparing this to Giere's colour analogy. One of the crucial points of that analogy for Giere is precisely that colours do not in any straightforward sense correspond to mind-independent facts; colour properties rather arise in the interaction of the visual system and the world. I will now turn to a form of perspectivism that takes this point seriously.

4.4. Relativist perspectivism

Prima facie, perspectivism is committed to relativism about truth. We have seen two approaches to perspectivism that attempt to avoid or at least downplay this commitment, but we have seen that both face difficulties accommodating incompatible models, and both have several other independent problems. I will now turn to relativistic perspectivism. It is this form of perspectivism that Chakravartty (2010) divides into two approaches:

(P1) We have knowledge of perspectival facts only, because non-perspectival facts are beyond our epistemic grasp.

(P2) We have knowledge of perspectival facts only, because there are no non-perspectival facts to be known.

A non-perspectival fact about a target system is thus a proposition that is true, independently of any particular perspective one may take with respect to it; it is true across perspectives. A perspectival fact is a proposition that is only true from, or within, or relative to, a given perspective (or limited set thereof).

Giere does not draw any distinction between moderate perspectivism, which avoids relativism, and the more radical perspectivism that embraces it. It is perhaps not surprising then that although Giere develops the similarity approach, which supports moderate perspectivism, he also explicitly endorses relativism:

For a perspectival realist, the strongest claims a scientist can legitimately make are of a qualified, conditional form: "According to this highly confirmed

theory (or reliable instrument), the world seems to be roughly such and such.” (2006: 5-6)

For a perspectivist, truth claims are always relative to a perspective. (2006: 81)

So Giere is sanguine about accepting relativism of some form, but this seems to be because he does not take truth to be of much importance in any case. Perhaps Giere would not see a distinction between what I am calling the moderate and radical approaches. Moderate perspectivism is already relativist, but the relativism is toothless since it is not judgments of truth but judgments of similarity that matter in the context of evaluating our models. However, I do not think that the consequences of relativism can be softened in this way. The problem is that it is a straightforward commitment of Giere’s own account that similarity relations are relative to perspective also. If truth is relative to perspective, and if the similarity claims are expressed in theoretical hypotheses which are straightforwardly true or false, then it is hard to see how appealing to similarity makes any difference with respect to relativism. Whether or not a model is similar to the world in a particular way turns out to be relative to perspective, in virtue of the relativity of truth. The point is that our access to the world is somehow indexed to or mediated by a perspective; whether it is truth relative to perspective or similarity relative to perspective does not seem to be especially important.

What exactly is relativism? I will follow Kusch (2017) in taking relativism to be committed to the following three theses:

Dependence (De): Beliefs have truth-value only relative to perspectives.

Plurality (Pl): There are many perspectives.

Non-Neutral Symmetry (NNS): There is no perspective-independent way to judge different perspectives.

I should note that Kusch is concerned with epistemic relativism, which endorses the relativity of *justification*. Our focus is alethic relativism, the relativity of *truth*. For epistemic relativism, (De) claims instead that a belief has an epistemic status, as justified or unjustified, only relative to perspectives. Obviously, other forms of relativism can be defined in terms of these three propositions, by substituting other factors for “truth value” in (De).

The important work here is done by (NNS). It may not be particularly controversial to state that beliefs have a truth-value only relative to perspectives, and that there are many perspectives, provided we have some absolute way of ranking alternative perspectives. We want to say that it is a fact – an objective, or mind-independent, or perspective-independent, or whatever, *fact* – that the Earth orbits the barycentre of the solar system. *From the Ptolemaic perspective*, this is false; but no matter, since we have good reason to think that the Ptolemaic perspective is wrong. If (NNS) is right, then our judgements of the epistemic status of different perspectives, and our attribution of truth and falsehood to different perspectives, are themselves perspectival.

It is crucial to distinguish relativism, particularly (NNS), from other theses with which it is sometimes conflated. Kusch distinguishes (NNS) from thesis of Equal Validity (EV):

(EV) All perspectives are equally correct, so there is no way to rank perspectives with respect to each other.

Many of the more colloquial arguments against relativism seem to be directed towards something like (EV): if we embrace relativism, the objector says, then we have no grounds for preferring science to voodoo, or vaccines to bloodletting, and so on. Relativism treats all theories and perspectives as on a par, and so anti-scientific, anti-rational, anti-intellectual. As Kusch argues, however, nothing of this sort follows from (NNS). (NNS) claims not that all perspectives are equal, but that there is no perspective-independent ranking of perspectives. (EV), by contrast, is a ranking: it ranks all perspectives as equal. Indeed, insofar as (EV) assumes a neutral standpoint from which a ranking is made, it would be a form of absolutism, not

relativism. In principle, a relativist might endorse (EV), where she ranks perspectives as equal from what she takes her own perspective to be, but nothing in relativism itself entails this commitment.

Additionally, relativism must be distinguished from:

Belief: P is true from your perspective just in case you believe that P.

(De), (Pl), and (NNS) do not entail Belief. It could be that, from your own perspective, many of your beliefs are false. In some cases, people are aware of the falsity of some of their own beliefs. A good example of this is the preface paradox: an author believes each individual proposition written in their book, but they also recognize their own fallibility, and so they expect that at least some of the propositions in the book are false. But even people who do not see any reason to lower their confidence in their own beliefs may be committed to propositions that are false per their own perspectives, even from the point of view of the relativist. For example, astronomers in the early 1800s, working in the Newtonian perspective, believed that there were only seven planets. This belief was false, per the Newtonian perspective; indeed, the Newtonian perspective played an important role in the construction of models that eventually led to the refutation of this belief. The discovery of irregularities in the orbit of Uranus, the hypothesis that Uranus was being perturbed by an unknown planet, and the calculation of the mass and position of the unknown planet, all took place within the Newtonian framework.

Finally, relativists need not accept:

Plenitude: For every proposition, there is some perspective from which that proposition is true.

Whether Plenitude is plausible perhaps depends on how perspectives are characterized. However, most perspectivists treat perspectives as models, or theories, or perhaps whole research traditions, and these may also include various methodological conventions, instrumental techniques, background assumptions, and so on. Recall Giere on maps:

I would like to say that the cultural background, the conventions for mapmaking, the designation of the region mapped, the specification of what features are mapped, and the degree of accuracy all determine a *perspective* from which the region is mapped. Every map reflects a perspective on the region mapped, a perspective built in by the mapmakers. (2006: 75)

Similarly, Massimi's definition of a "scientific perspective" as:

the actual – historically and intellectually situated – scientific practice of any real scientific community at any given historical time. I understand scientific practice broadly to include: (i) the body of scientific knowledge claims advanced by the scientific community at the time; (ii) the experimental, theoretical, and technological resources available to the scientific community at the time to reliably make those scientific knowledge claims; and (iii) second-order (methodological-epistemic) claims that can justify the scientific knowledge claims so advanced. (Massimi 2016a: 2)

Given these characterizations, there is no reason to expect that just any proposition will turn out to be true, providing we choose the right perspective. There is surely no perspective from which it is true that the Earth is wholly populated with fairies, for example.

Having specified the commitments of relativism, we can now consider how this form of perspectivism deals with the incompatible models argument. We saw that the basic issue, from a realist point of view, is that incompatible models undermine the success-to-truth inference. We are presented with several models that exhibit the relevant kind of success but which, it is argued, cannot all be true, in virtue of incompatibility. Relativist perspectivism offers a straightforward solution to this, because it attempts to provide an account on which these successful models are all taken as *true*. The general idea is simple enough. Suppose we take it that there are "perspectival facts"; facts which are revealed only from a particular perspective, or propositions that are true only relative to a given perspective. Now suppose that

model M1 supports claim X about some system, and M2 supports claim not-X. The standard view is that both X and not-X cannot both be true, so at least one of these models must be wrong. The perspectivist relativizes the truth to perspectives: X is true-from-P; not-X is true-from-P*; and then makes the additional claim that there is no “non-perspectival” view from which perspectives can be ranked. More broadly, we evaluate the models as a whole from a given perspective. M1 is “true”-from-P; M2 is “true”-from-P* (the scarequotes are used here because models are generally thought of as nonlinguistic, and so do not have truth values; I use the term here just to designate some relation of rightness, such as similarity). Similarly, with respect to the diachronic argument, antirealists point out theories of the past that were successful in ways that realists take as indicative of truth, but that have since been displaced by incompatible superior theories. On a perspectivist account, these past theories may well count as true, but true relative to the perspectives prevailing at the time.

There is a straightforward sense in which relativist perspectivism allows for the defence of a limited kind of realism. Of course, traditionally, realism and relativism have been seen as opposed; and we might well simply define “realism” such that relativism is incompatible with it. I would make two points here. First, the whole point of perspectivism is offer a *via media*, a position that accommodates the intuitions driving both standard realist and antirealist positions and carves a path between them. The realism of perspectivism is explicitly a limited realism. Second, standard definitions of realism do not usually rule out relativism – at least, not obviously so. Scientific realism, broadly speaking, is committed to three theses:

(SR1) There is a mind-independent world.

(SR2) Scientific theories yield literal descriptions of the world.

(SR3) We are justified in believing that some of these claims are true.

These claims would need to be preserved in a perspectival *realism*. As the position has been elaborated here, all these claims can be accepted by a perspectivist. Indeed, the perspectivist can preserve a fourth, even stronger claim, namely:

(SR4) The success of science is explained by the truth of our best theories and models.

This explanatory claim, which underlies the success-to-truth inference, is supported by the perspectivist's response to the incompatible models argument. Under standard forms of realism, incompatible models undermine the success-to-truth inference as examples of success without truth. The perspectivist can resist this.

What a perspective realist will add to these four theses is that truth is relative to perspective. True claims can be made within perspectives (De), various perspectives can be and have been constructed (Pi), and there is no non-perspectival way of evaluating any given perspective, or no non-perspectival truth for any given perspective (NNS). It is this explicitly relativist component that distinguishes perspectival realism from the many other types of realism. Perspectival truth is then simply truth relative to perspective; and the defining feature of perspectivism is that all truth, at least within the sciences, is perspectival truth. The goal of the rest of this thesis is to elaborate on what exactly this amounts to.

Relativist perspectivism connects very straightforwardly with the incompatible models argument. It is also clearly distinguished from standard realist and instrumentalist approaches. As such, relativist perspectivism at least has the potential to offer a genuine *via media*. But to grant this much is, of course, a far cry from showing that the argument from incompatible models successfully establishes relativist perspectivism. First, one might raise objections to relativism, and attempt to show that this is incoherent – if the relativism is incoherent, then perspectivism is not a genuine alternative to realism and instrumentalism, after all. Second, alternative approaches to incompatible models might be proposed. If incompatible models can be accommodated in other positions, then the grounds for perspectivism are undermined. We will take these points one at a time.

4.5. The split strategy

There is an important qualification in the penultimate paragraph of the previous section. All truth, at least within the sciences, is perspectival truth. This suggests that in addition to the perspectival truths generated by science, there are also non-perspectival truths. This is indeed a line that some relativist perspectivists have

taken. This is understandable. Relativism is extremely controversial, even granting the clarifications made in the previous section. In particular, there have long been challenges to the coherence of relativism. Two arguments are standard (see Kukla 2000 for discussion):

(a) Self-refutation

Relativism claims that all truths are relative to some perspective, and that there is no non-perspectival way to judge different perspectives. So, is relativism true or false? The relativist faces a dilemma. Either relativism is absolutely true, or it is only relatively true. If it is absolutely true, then there is an absolute truth, and we have abandoned relativism. If it is only relatively true, then it is false for absolutists, hence relativism is, again, absolutely false. The relativist has conceded that the absolutist is correct, from the absolutist's own perspective.

(b) Infinite regress

The relativist wishes to claim that all truths are relative to some perspective. But this claim is itself also true only relative to some perspective. We have:

(R) All truths are relative.

(R1) (R) is true relative to A.

(R2) (R1) is true relative to A*.

(R3) (R2) is true relative to A**.

And so on.

We have an infinite regress of perspectives. But presumably there could not be an infinite sequence of perspectives. Perspectives are constructed by humans – as we have seen, they consist of sets of knowledge claims, methodological conventions, instruments and technological practices, and so on; clearly, we do not have an infinite number of these.

A slightly different way to put the regress problem is to say that relativism cannot even be expressed. There is no proposition that captures what the relativist wants to say (this is how Kukla frames the regress problem; see Kukla 2000: 132). Suppose for example that the relativist asserts some proposition, say, “the sky is blue”. Now,

the relativist cannot accept this without qualification, as this does not express her actual state of opinion. Instead, if she is speaking precisely, she must say something like:

“the sky is blue” is true relative to A.

And presumably she would then add something like, “I accept/endorse A”. The problem is that in saying this much, she has again failed to express her actual opinion, which would need:

““the sky is blue” is true relative to A” is true relative to A*...

Here is the regress. This is particularly troubling to the relativist, because it implies that relativism cannot even be expressed. The relativist can never state any proposition that captures what she wants to say.

Both problems arise for a “global” relativism, the position that all truths are relative. So, an obvious way to make relativism more plausible is to restrict its scope.

Whatever else we might say about moral relativism, for example, it does not seem to be self-refuting, or to face a vicious regress, because the proposition, “moral truths are relative to moral perspectives” is stated as an absolute, non-relative truth.

Perspectivists can make the same move, holding that while science delivers perspectival truths, there are also non-perspectival truths. This seems to be Giere’s position; in an article on Kuhn, Giere attempts to distinguish his position from that of “global perspectivism”, saying that “the perspectivism involved is not global, but confined to scientific knowledge, so a *scientific* perspectivism. The presupposed conceptual scheme is the property of a scientific community” (Giere 2012).

Unfortunately, Giere does not elaborate on this point, but it seems that he accepts perspectivism only with respect to the sciences. This is to say: within the sciences, truth is relative to a perspective (De); there are many scientific perspectives (PI); and the sciences do not provide any perspective-independent means of judging different perspectives (NNS). Of course, this is an important qualification, because it raises the possibility that there may be a perspective-independent way of judging scientific perspectives, it’s just that we would have to look outside the sciences proper for this.

This is what I am calling the “split strategy”, because it involves proposing that there are two different kinds of truth: perspectival truth and non-perspectival truth. This same approach is suggested by Elgin’s “constructive nominalism” (Elgin 2019). Elgin defends a form of relativist perspectivism, and as with many other perspectivists, the central motivation is the multiplicity intuition and the problem of incompatible models. Elgin’s solution to the problem is to say that what makes for a good theory is determined by factors other than truth, such as empirical adequacy, simplicity, robustness, explanatory power, and so on, and then truth is constituted within the theory. So, if the wave theory of light is assumed to meet all the relevant theoretical virtues, then “light consists of waves” is true, though this is shorthand for “as modulated through theory T, light consists of waves” (Elgin 2019: 532-533). What makes the wave theory of light a good theory is not that it is true, but that it exhibits a host of theoretical virtues. A similar approach is Chang’s (2017; 2018) “pragmatist coherence” view. Like Elgin, Chang takes our judgment of theories to be based on pragmatic factors, and truth to be constituted within the theory. As Chang says: “we regard as real and true the entities and statements that underpin coherent activities.” The primary difference between Chang and Elgin seems to be that Chang is more permissive about just what counts as a good theory. For Chang, phlogiston chemistry was a coherent activity, at least for some chemists, and so even phlogiston is real.

Perspectivism here is limited to theoretical claims. Truths concerning the phenomena need not be taken as constituted within any particular perspective. We do have straightforward, non-perspectival access to what we observe, or at least to what we experience. So “light consists of waves” is a perspectival truth, a truth constituted within a given theory of light, but “there is an Arago spot in the shadow” is simply true.

However, there are problems for the split strategy. First, there is the question of its motivation. Science is our best means for generating knowledge of the world, and one reason for this is that it seems to achieve objectivity, in some sense at least (Reiss and Sprenger 2020). What makes scientific perspectivism philosophically interesting is precisely that, at least *prima facie*, the sciences seem to present the most challenging domain in which to take a perspectivist position. Science, when it

works well, provides a view of the world divorced from the biases and idiosyncrasies of particular investigators or communities. If scientific facts are perspectival, what facts exactly count as non-perspectival? No doubt most perspectivists would not want to claim that scientific knowledge is in any sense *defective* in comparison to other forms of knowledge (perceptual knowledge, philosophical knowledge, or whatever), but this seems to be the view of those who adopt the split strategy.

Perhaps the obvious option is to try to make a distinction between theoretical and observational claims, so that direct observation statements are non-perspectival. Of course, this distinction has a long history of difficulties. But more importantly, as we have already seen, Giere devotes an entire chapter to the perspectival nature of perception, and he explicitly denies that direct perception is on a different footing from other ways of forming knowledge. Indeed, colour perception provides the paradigm for a perspectivist analysis. Even if a strict distinction between theoretical and observational claims could be drawn, the motivation for applying perspectivism only to the former is unclear. The best case for perspectivism is perception; and we frequently use instruments and general theoretical assumptions to correct beliefs formed from perception.

The most serious problem for the split strategy is why exactly we should accept this dual notion of truth. Suppose that “light consists of waves” is only perspectivally true, while a claim such as “there is an Arago spot in the shadow” is non-perspectivally true. The natural response to this, it seems to me, is to say that “perspectival truth” is just not really a kind of truth, as “truth” is commonly understood. The perspectivist is engaging in semantic sleight-of-hand. It is our everyday notion of truth that is found in the truth of the statement “there is an Arago spot in the shadow”, and in *this* sense, theoretical claims such as “light consists of waves” are not true – or at least, there is no reason to think they are true. There is, of course, nothing to stop us from introducing some alternative way of thinking about truth, which allows us to say that various theoretical claims are true-within-the-theory, but from the point of view of a standard realist, or even anybody just on the fence, talk of theoretical and observational claims as both being *true* will surely seem like a bait-and-switch. What exactly is the difference between saying that P is perspectivally true, and saying that

P is merely pragmatically useful (as the constructive empiricist might say for some theoretical claim)?

Compare, for example, the proposal that to understand claims about fiction, we need a concept of fictional truth, where fictional truth is distinguished from ordinary truth in that claims which are fictionally true must be understood as containing an unstated reference to a particular fictional work: "P is fictionally true" = "according to fiction X, P is true." This will allow us to claim that various propositions are fictionally true, e.g. it is fictionally true that Sherlock Holmes lives on 221B Baker Street, where the context makes it clear that we are talking about the works written by Arthur Conan Doyle. There may well be good reasons for introducing this concept of fictional truth. But nobody would think that this supports a realist reading of the works of Arthur Conan Doyle. To say that P is only fictionally true is to say that it need not be true as a literal description of the real world. The same problem arises for perspectival truth, when contrasted with non-perspectival truth. The global perspectivist, by contrast, can hold that in talking about truth, we were talking about perspectival truth all along. For the global perspectivist, "perspectival truth" is proposed an analysis of our everyday concept of truth.

4.6. The coherence of relativism

If the split strategy doesn't work, then the perspectivist must accept a global relativism. We have noted two challenges to the coherence of relativism. However, both rest on a mischaracterization of the position.

(a) Self-refutation

First, even if we take this argument at face value, it is not clear that relativism is self-refuting. The claim is that if the relativist takes relativism itself to be only relatively true, then she must accept that, from the absolutist's perspective, relativism is false. But to assume that this self-refuting is, as Giere points out (2006: 81), straightforwardly question-begging, since it assumes that truth is to be analysed in absolutist terms (see also Bloor 1976). To say that relativism is false, from some other person's perspective, is not itself to deny relativism, given the relativist understanding of truth. At no point then does the relativist assert the contradiction

that relativism is true and that relativism is false. She asserts that relativism is true, and then takes this, like all other truths, to be relative to her perspective. She can acknowledge the existence of other perspective without asserting their content.

The most that can be said here is that the relativist has *failed to deny* absolutism. The relativist must grant that, from the absolutist's perspective, relativism is false. But there are two reasons why this is not compelling as a response to relativism. First, grant the assumption that relativism cannot deny absolutism. This is not a conclusive objection. To see why, consider trivialism. Trivialism is the view that all propositions are true. It follows, of course, also that no propositions are true, since trivialists assign "true" to the proposition "no propositions are true", and that only some propositions are true, since trivialists assign "true" to the proposition "some propositions are true". Trivialism is a bizarre position. It's incoherent. It's impossible to believe it. Yet some philosophers have defended it (Kabay 2010 defends trivialism and argues that there have been defences of trivialism historically). The relevant point in this context is that in the same way that it's not possible for the relativist to deny absolutism, it's not possible for anybody to deny trivialism (cf. Kabay 2010). Since the trivialist claims that all propositions are true, she will agree with everything you say. Whenever you attempt to state an alternative to trivialism, you are simply asserting part of the content of trivialism. You say: "trivialism is false" – well, the trivialist agrees that trivialism is false. "Trivialism is false" is part of the content of trivialism. I suggest that just as our inability to deny trivialism is not an automatic win for the trivialist, the relativist's failure to deny absolutism is not an automatic win for the absolutist.

In any case, putting this point aside, the self-refutation objection fails for a simpler reason. The relativist *can* deny absolutism. This is because she can argue that relativism is true *from a perspective that all interlocutors accept*. She can argue that the absolutist is mistaken about the commitments of their own perspective. In practice, this is how arguments for relativism proceed. The relativist will cite premises and uses rules of inference that, she hopes, everybody accepts, and then tries to show how these support relativism. This is the point of the incompatible models argument. The relativist points to phenomena that we all recognize, and then tries to build a case for relativism on these shared assumptions. The self-refutation

objection assumes either Belief (P is true from your perspective just in case you believe that P) or Plenitude (for every proposition, there is some perspective from which that proposition is true): that absolutism is true from your perspective provided you believe absolutism, or that there is some perspective from which absolutism is true. We have seen that relativism need not be committed to either of these theses. It may be that from the absolutist's perspective, absolutism is false, or even that from every perspective, absolutism is false.

(b) Infinite regress

This objection assumes that when the relativist asserts some proposition P, this must be understood as abbreviating the claim, "P is true relative to perspective A". Now, the relativist will, of course, assert that P is true relative to perspective A. But this need not entail that, when she asserts P, she thereby asserts that P is true relative to A. The point here is that "P" and "P is true relative to A" are different claims. The relativist will assert both, but the former need not be taken as abbreviating the latter. The relativist's position is adequately captured by (R):

(R) All truths are relative.

Of course, the relativist will also endorse (R1):

(R1) (R) is true relative to A.

But (R) already expresses the relativist's opinion. (R) need not be taken as abbreviating (R1).

Consider again an everyday claim such as "the sky is blue". The argument is that this does not express the relativist's actual opinion, which would require the qualification:

"the sky is blue" is true relative to A.

The relativist has no reason to accept this. She will assert that the sky is blue. Her analysis of truth is then a separate claim. It is not part of the content of her original

assertion. Similarly, when the absolutist says “the sky is blue”, we do not take this to be abbreviating:

“the sky is blue” is true absolutely.

The regress argument assumes that when the relativist makes any claim, she must add the qualification that it is true relative to some perspective in order to avoid affirming anything absolutely. But this is simply to assume that truth must be understood in absolutist terms. You would only accept that this qualification must always be stated if you already accept an absolutist account of truth. It is unsurprising then that this same argument can be turned against the absolutist. When the absolutist says, “the sky is blue”, the relativist can object that she has failed to state her actual position, because so far, she has only said that it is relatively true that the sky is blue. The absolutist needs to add the qualification:

“the sky is blue” is true absolutely.

But now this generates a regress, because she must also add that this sentence is true absolutely. For every proposition P that the absolutist affirms, she needs the higher-level proposition ““P” is true absolutely.”

If we accept Tarski’s T-schema then it is plausible to say that when a person asserts “P”, they thereby assert ““P” is true”. The former is, perhaps, an abbreviation of the latter; the assertion of P is an assertion of the truth of P. But it strikes me as extremely implausible to suppose that when a person asserts “P”, they thereby assert either ““P” is true relative to some perspective” or ““P” is absolutely true”. The debate between relativism and absolutism is a question for philosophical analysis. Is truth to be understood in relativist or absolutist terms? Our position on this need not be part of the content of any particular proposition we assert.

It turns out that the split strategy, in addition to facing several difficulties, is not well-motivated in the first place, at least insofar as the goal is to avoid the traditional objections to relativism.

Throughout this chapter, I have used the argument from incompatible models to explore various forms of perspectivism. We can classify the various perspectivisms according to how they answer three questions: Truth or alternative representational relation? If truth, relativist or non-relativist? If relativist, global relativism or local relativism? I have argued that only a relativist perspectivism do the job of carving a *via media* between traditional realism and antirealism, and that the best prospects lie in a global relativism. Our final question is whether the argument from incompatible models provides a convincing case for this form of perspectivism.

4.7. *Alternative accounts of incompatible models*

Despite its prominence in the perspectivist literature, the argument from incompatible models has received a great deal of criticism. Ultimately, I do not think that it provides a plausible foundation for perspectivism. The standard criticism of the argument from incompatible models is that there are various other, less controversial strategies available for accounting for incompatible models, which undermines the use of incompatible models to support perspectivism. This is discussed by Chakravartty (2017) and Morrison (2015). These strategies are available to straightforward realists and antirealists. Call this the “alternative strategies response”.

(1) Resolve the incompatibility by taking the models to be representing different properties. In this case, it turns out that the models are not actually incompatible. Morrison (2015: 168) argues that most cases where perspectivism seems to provide an appropriate framework can be better understood in this way. Consider the proposal that scientists use incompatible models of fluid flow. In fact, scientists are simply modelling “different types of behaviour in different parts of the fluid.” If our purpose is to model fluid flow, then we use *these* models; if our purpose is to understand diffusion, then we use *these* models (cf. Giere 2009). A “perspective” on the fluid appears to involve nothing more than the selection of properties of interest, and for different purposes, we may be interested in different properties. So naturally, we should expect to generate different models. Additionally, it is no surprise that the idealizations we use in these models may also be different, so we end up with a variety of seemingly incompatible models. This is no obstacle to non-perspectival

knowledge. We are simply using a variety of models to build up an overall picture of the behaviour of the fluid.

(2) In some cases, we have some underlying theory that explains why the models give rise to incompatible representations. Maps provide an excellent example of this: we can generate a host of incompatible map projections when mapping the Earth's surface on a 2D plane. In this case, we can "step outside" any particular projection: knowing how the projection was constructed, we can work out how the surface would look as a globe; and knowing the shapes and sizes of the countries on the globe, we can understand this is transformed into a 2D projection:

the variability of the resulting appearances may be explained in terms of *actual* shapes and sizes and the mathematics of projections. In just the same way, the appearances of a straw as straight in the air and bent in a glass of water are explainable in terms of the actual shape of the straw combined with theories of optics and the human visual system (Chakravartty 2017: 174)

The same can be done for many cases of scientific models: for instance, we can account for the success of Newtonian models of the solar system in terms of what we take to be the facts (or something closer to the facts) as described by general relativity. For various purposes, the distortions introduced by Newtonian models are simply irrelevant. This is of course a statement of the general problem of Escape from Perspective, applied to the specific case of idealized models. Moreover, the underlying theory need not be anything particularly sophisticated. When scientists construct models, they regularly employ idealizations. In doing so, they know that they are distorting certain features of the world. But if we can specify that some aspect of a model is a distortion, an idealization, etc., then we must know something about the corresponding state of affairs of the target system, in a way that is independent of the model in question. It is because we know, independently of any specific biological model, that biological populations are not actually infinite that we can say that the infinite population in a biological model is an idealization.

(3) In some cases of incompatible models, one model is clearly epistemically superior to another. Here the incompatibility is straightforwardly resolved: we have good reasons for treating only one of the models as a correct description of reality. We might see this as a special case of (2); Chakravartty cites Newton and Einstein's gravitational theory here. It is because Einstein's theory is epistemically superior that we take it to provide a description of the underlying facts, and this allows the application of (2) to the Newtonian models. Along the same lines, Giere himself takes it that such comparative judgements of perspectives are important, partly to save the realist aspect of perspectivism. When we model water:

We can understand how large numbers of small molecules might behave like a continuous fluid. We cannot understand the phenomenon of diffusion from a fluid mechanics perspective. That asymmetry is all that a perspectival realism requires. (Giere 2009: 222)

Thus, we can understand the fluid mechanics perspective from the point of view of the molecular perspective, but not vice versa. This gives us reason for taking the molecular perspective to be epistemically superior. In this case, the models themselves allow for a realist interpretation, despite their incompatibility. The key point is that we don't simply have a collection of separate models; the relations between our available models are often rather complex and allow for comparative judgments of models. But now it appears that we have, again, another route to Escape from Perspective. Water in fact simply consists of innumerable small molecules. That's just how water is constituted, and we can understand its fluid behaviour in these terms. When then is the point of making the additional claim that this is a perspectival fact or a perspectival truth? Giere may have saved the realism, but he seems to have lost the grounds for perspectivism. Where one model can be selected as superior, we can treat that as the model that gets the underlying ontology right, while other, incompatible models can be treated as useful tools. Knowing what water is – knowing that it is not actually continuous but composed of molecules – we can see why the more accurate models would be practically inapplicable in this context, perhaps because the calculations involved would be far too complex, or something along those lines. So, we develop idealized models as instruments.

Strategies (1), (2), and (3) support a standard realist reading of models via Escape from Perspective. Of course, there may be much more to be said about these cases. We might wonder about the precise relations between, say, the superior and the inferior model in (3): is this to be understood in terms of reduction, or elimination, or something else? Similarly, we might argue that there is still room for truth within the inferior model. No doubt, the inferior model must describe some aspects of the target system correctly, if only aspects of its observable behaviour. But there is no incompatibility here, because it is only specific propositions drawn from the model that we are taking to be true, which may be consistent with what we think we know from the superior model.

(4) We may simply withhold belief. The question facing Newtonian mechanics of whether the universe is in motion or at rest is a case of this. Given the evidence available at the time, no belief should be held. Now the problem is that in cases of genuinely incompatible models, where there is nothing available to break the tie between different models, it would appear that (4) is the reasonable response. We simply do not yet know enough to make a judgement. Morrison's argues that with nuclear models, there are a variety of models that treat the underlying structure of atoms differently, with nothing to decide between them. In this case, there is nothing to favour perspectivism over straightforward instrumentalism. Indeed, it's notable that appears to incompatible models have already been used as arguments for antirealism; see Morrison's own earlier work (1990). The argument here is straightforward: If two models are incompatible, we know that at least one of them must be wrong; if both are about equally successful, our justification for believing that either model provides an accurate representation is undermined. As a response to perspectivism, this would be question-begging, since the perspectivist denies that the incompatibility entails that at least one model is wrong. My point is simply that the existence of incompatible models gives us no grounds on which to decide between antirealism and perspectivism.

There may be other strategies for dealing with incompatible models, but this gives us enough to see how standard realists and antirealists should proceed. For antirealists, the situation is simple. They can grant that scientists regularly develop incompatible

models but treat this as grounds for antirealism. As for realists, they will apply strategies (1), (2), and (3) where they can, thus reducing the number of successful incompatible models. For the residual cases not susceptible to analysis in these terms, the realist can just apply strategy (4) and accept instrumentalism. After all, there is an enormous variety of models, and an enormous variety of purposes to which models are put, and nobody is a realist about all of them. Traditional scientific realism is restricted to those theories and models that exhibit an appropriate degree of success, and it is reasonable to hold that one aspect of success is coherence and consistency. So, there is nothing *ad hoc* in taking an instrumentalist view of genuinely incompatible models. They may exhibit a great deal of predictive success, but this is undermined precisely by the fact that they are incompatible.

This points to a more general problem. It is not obvious how to argue for any alternative to traditional realism by citing specific features of scientific practice. Realists are already instrumentalists about most of the theories and models that have been developed. Whatever a perspectivist points to as providing support for perspectivism, it will be open to the realist to just take an instrumentalist line on it, without generalizing this to all the products of science.

Is the alternative strategies response compelling? This depends on how exactly the perspectivist is conceiving of the incompatible models argument. If the claim is that perspectivism is the only way to understand the proliferation of incompatible models, then the existence of alternative strategies seriously undermines the perspectivist's argument. But this is not necessarily how the perspectivist must proceed. I noted, when presenting the incompatible models argument, that the difficulty this raises for the standard realist is that it undermines the success-to-truth inference. Granted, *if* we accept realism, then we can account for incompatible models by applying these alternative strategies. But the question is why we should accept realism in the first place. Notice that other than strategy (1), all these alternative strategies involve making comparative judgments. We must take some of the models to be epistemically inferior, and so take the target system to be correctly described by other models, as in strategies (2) and (3), or just be agnostic about the relevant properties of the target system, as in strategy (4). The problem here is that these inferior models are still successful, and they may exhibit the relevant kinds of

success that, in other contexts, the realist takes to be truth-tracking. Recall Kitcher's success-to-truth rule:

S plays a crucial role in a systematic practice of fine-grained prediction and intervention.

So, S is approximately true.

When we identify models that exhibit the relevant kinds of success, but that are incompatible, we have apparent counterexamples to the success-to-truth inference. This is the case even if we can apply these alternative strategies. Take models M1 and M2, and suppose we can understand M2 from the point of view of M1, but not vice versa. Thus, we take M1 to provide a correct description of the underlying properties of the target system. M2 is understood instrumentally, as a useful tool we can apply that allows for more simple calculations, or something along those lines. Now if M2 is used to derive successful novel predictions, or is applied in explanations of phenomena, or is used in engineering contexts, etc., then M2 is a counterexample to the success-to-truth inference. The perspectivist offers a way out of this problem. If truth is to be understood as truth-relative-to-perspective, then both M1 and M2 can be taken as true, or at least the propositions we infer from M1 and M2 can be taken as true, without incompatibility. The apparent counterexample is removed.

Of course, this still does not show why perspectivism should be preferred to outright antirealism, who can also offer a straightforward analysis of incompatible models. If the only way to save a form of realism is to embrace this kind of relativism, many might find antirealism preferable. Moreover, there is a worry that the perspectivist's response will be too general. There are numerous apparent counterexamples to the success-to-truth inference, and in many cases, it seems implausible to take them as true in any sense, even if we understand truth in perspectivist terms. Is the perspectivist committed to the existence of phlogiston, caloric, the luminiferous aether, the celestial spheres, etc.? All of these were involved, in one way or another, in theories and models that exhibited a striking degree of empirical success. Even the perspectivist will sometimes want to make use of strategies (1) to (4) to deal with such cases. But then even the perspectivist will face apparent counterexamples to

the success-to-truth inference. Of course, the perspectivist can also respond to such cases by arguing that the models in question were not successful, that they do not meet the required standards of success for the success-to-truth inference to be applicable. The same move, however, can in principle be made by the realist in response to any examples that the perspectivist cites.

Any perspectivist who uses the argument from incompatible models must perform a delicate balancing act. She must point to cases of incompatible models, where both models meet the standards of success required for the success-to-truth inference. For many incompatible models, the perspectivist will want to say that at least one of the models is straightforwardly false, that strategies (1) to (4) can be applied. If the success-to-truth inference is to remain plausible, these cases need to be minimized, so the perspectivist needs to argue that such models do not actually exhibit the relevant success. But she also needs to provide a convincing reason why the realist should not simply use the same response against the examples that she cites. In other words, suppose we have two incompatible models of one system, M1 and M2, and two incompatible models of another, M3 and M4. These models are all successful in some ways. Suppose that M1 is a caloric model. The perspectivist says that M1 is simply false in this respect. She will then argue that M1 does not exhibit the kind of success required to infer truth, hence it does not pose a challenge to the success-to-truth inference. But both M3 and M4, according to the perspectivist, do meet this standard, so to maintain the success-to-truth inference, we must take each as true relative to different perspectives. Naturally, the traditional realist will reject this. The fact that M3 and M4 are incompatible is good reason, from a realist point of view, for thinking that it's not the case that both meet the relevant standard of success. There is enough debate about just what counts as "success", and under what conditions success is indicative of truth, that this response is plausible.

At the beginning of this chapter, I distinguished two questions. First, is perspectivism a genuine alternative to standard realism and antirealism? Second, is there any reason to accept perspectivism? I conclude that although moderate forms of perspectivism turn out to be indistinguishable from standard realism, there is a distinctive, relativist form of perspectivism available. This is an important point in the perspectivist's favour. As we have seen, critics of perspectivism have accused it of

collapsing into either realism or antirealism. However, the argument from incompatible models does not provide a convincing reason to accept this form of perspectivism, given the availability of alternative strategies for understanding incompatible models.

Chapter 5

Failure of Fit

5.1. Language and reality

I have argued in the previous chapter that the incompatible models argument does not provide a good justification for accepting perspectivism. However, this argument is only one way of elaborating on the underlying “multiplicity intuition” that motivates perspectivism: the idea that truth is revealed from multiple different viewpoints. There is a second argument, or perhaps family of arguments, which builds on this intuition, and which is sometimes presented in connection with the argument from incompatible models, which I will call the argument from failure of fit; this is argument is made most notably by Giere and Teller. The basic claim of this argument is that our representations fail to hook up to the world in the kind of way required by traditional realism, and that only a perspectival realism can accommodate this. One statement of this argument is given by Kuhn:

Evaluation of a statement’s truth values is...an activity that can be conducted only with a lexicon already in place, and its outcome depends on that lexicon. If, as standard forms of realism suppose, a statement’s being true or false depends simply on whether or not it corresponds to the real world – independent of time, language, and culture – then the world itself must be somehow lexicon-dependent. Whatever form that takes, it poses problems for a realist perspective, problems that I take to be both genuine and urgent. (2000: 77).

All scientific theories involve taxonomies or lexicons, which order the world in different ways. The traditional realist appears to be committed to the view that, in some sense, our theories speak the language of the world. To give an objective account of some object is, as Reynolds (2018: 194) puts it, to give “a description in the object’s own terms or language.” This idea has been noted by Giere in a couple of recent articles (2012; 2016). Recall that the challenge for the perspectivist is to make sense of the notion of perspective-dependent truths or perspective-dependent facts. But we might wonder whether this challenge is reasonably motivated in the

first place. What sense is there to the notion that truth could be perspective-independent? Assume that truth consists in some sort of relation between a proposition and the world. A correspondence theorist, for instance, will say that a proposition is true just in case the proposition “corresponds” to the world, just in case it “matches” the world in an appropriate way. As stated, this is just a metaphor, and this metaphor can be developed in various ways. But it captures an important intuition about the nature of truth; and plausibly, any acceptable theory of truth needs to accommodate this intuition. To say that truth involves some sort of “matching” relation between a statement and the world is simply to express the commonsense point captured in Tarski’s material adequacy condition:

“P” is true iff P.

e.g. “snow is white” is true iff snow is white.

or “schnee ist weiss” is true iff snow is white.

On the right-hand side is a proposition about the world. We need not understand “the world” here in any metaphysically loaded sense. Surely, however, any kind of realist will suppose that we do have access to the external world, in some sense of the term, and that we can say true things about it. Now, truth presupposes a language, in which one can form statements that have a truth-value, and this involves some set of concepts. There can be no truth independent of human conceptual schemes. If truth involves a relation between statements and the world of the kind that drives the correspondence theory – or perhaps, of the kind that must be respected by any plausible theory of truth – then the world itself must have something like our conceptual structure somehow built into it. That is, to suppose that there can be perspective-independent truths, we would have to assume that the world itself has a kind of linguistic or conceptual structure that is inherited by our representational systems. Here the perspectivist objection enters: there is reason to suppose that there is some kind of mismatch between the world and our representational capacities. Thus, Giere says of traditional realism that it:

incorporate[s] a metaphysical account of truth according to which there are truths about the world that exist independently of human existence ... This account requires that the world itself contain something like “facts” that

mirror the linguistic structure of statements describing them. But this leaves us with the problem of understanding how the world itself could have independently acquired a structure corresponding to human languages, which are humanly created artifacts. (2015: 189)

Giere gives another statement of the argument in a slightly later article:

Taking some hints from Nietzsche, one should begin by insisting that truth is a relation between a statement and the world. So any talk of truth presupposes a language. Truths independent of human conceptual schemes would have to be understood as related to the world's own language, so to speak. That idea, as Nietzsche points out, made sense assuming the Christian God who spoke the world into existence. The world's language is God's language. On Nietzsche's telling, Enlightenment philosophers killed God but failed to realize that meant eliminating God's truths as well. Contemporary philosophers, including most scientific realists, are making the same mistake. One would have to presume that discovering truths about the world required first coming up with the right set of concepts, in effect, discovering the language that a Christian God would have used. Given that human languages are historically evolving representational systems, and that the world is incredibly complex, that seems most unlikely even if not logically impossible. (2016: 140)

Giere makes the same point about laws (2006: 70). Originally, Giere says, talk of laws was drawn from Christian theology; laws of nature were thought of as being literal laws, laid down by a lawgiver. The world is governed by laws of nature, in that it follows God's commands. Here, again, is the notion of the world being "spoken into existence" by God. Over time, explicit theological notions dropped out of our understanding of science, but scientists and philosophers continued to think of phenomena as governed by laws. (See also Garber 2016.)

Although Kuhn and Giere frame this point in terms of lexicons, or in terms of languages, this should perhaps be taken as a metaphor. As we have seen, one of the primary representational tools in the sciences are models, and models are not

generally thought of as linguistic structures. The general idea of failure of fit is that there is a mismatch between our representations and the structure or nature of the world itself. Additionally, nothing in the general failure of fit argument requires that our representations, or the lexicons of our theories, are as rigid as the lexicons of natural languages. Perhaps Kuhn saw things this way: for Kuhn, a theoretical lexicon may be relatively stable and difficult to shift, until anomalies in the theory lead to a revolutionary paradigm change. All scientists working in a particular domain during non-revolutionary normal science would be expected to largely share the same lexicon. Giere, however, rejects this aspect of the Kuhnian picture (Giere 2006: 82); for Giere, we can often shift between competing perspectives quite easily, and we can often understand the lexicon of one perspective from the point of view of another. However, Giere still accepts the apparent relativist consequence of this position:

For a perspectivist, truth claims are always relative to a perspective. This is not so radical a view as it might sound. It was long a doctrine within Logical Empiricism, and analytic philosophy more generally, that scientific claims are always relative to a language. First one chooses a language, then one makes claims that may be judged true or false. And the choice of a language is pragmatic, not itself a matter of truth or falsity. (2006: 81)

It is worth noting that in several ways, this argument has better prospects for justifying perspectivism than the argument from incompatible models. First, rejecting, or at least remaining agnostic about, the lexicon-dependence of the world, as Kuhn puts it, is one way to reject, or remain agnostic about, perspective-independent facts. So, this argument can straightforwardly ground a metaphysical, not just an epistemic, form of perspectivism, which is what we are examining here. If this argument succeeds, then there are only perspective-dependent facts. By contrast, the argument from incompatible models can be read as drawing only an epistemic conclusion. Of course, if there are only perspective-dependent facts, then an epistemic perspectivism also follows trivially, since we would only have access to perspective-dependent facts.

Second, more significantly, recall that the central problem for the argument from incompatible models was that it is difficult to use it to justify perspectivism over other approaches, because other interpretations of incompatible models are always available. The incompatible models argument therefore does not provide the resources to avoid Chakravartty's dilemma. The problem is that very often, we can "step outside" a given model, and thus apply Escape from Perspective; whereas in situations where we have no means to step outside the available models, we can conclude that all models are false or that only one model is accurate, but remain agnostic about which it is. Realists are not required to be realists about every model that scientists use.

For the argument from failure of fit, the situation is different. The point of this argument is that there is, in principle, reason to expect that we are unable ever to "step outside" our perspective in the relevant sense. There are, in principle, limits to the extent to which the Escape from Perspective move can be applied. Of course, this still leaves us with the option of adopting an instrumentalist interpretation. But if the argument from failure of fit is successful, then we would have to adopt instrumentalism across the board. The failure of fit applies to all domains, including claims about perception: Teller (2018a), for example, argues that problems of reference failure apply equally to theoretical posits and to everyday observables. There may be an epistemically relevant difference between observables and unobservables – I have already made a proposal for one way of drawing this line, in the chapter on instrumental perspectivism – but the failure of fit argument provides an additional challenge to traditional realism which cuts across this dichotomy. So adopting an instrumentalist approach in this case would commit us to radical skepticism, at least if we assume a realist view of what knowledge is supposed to consist in, i.e. a kind of mirroring/correspondence with the mind-independent world. This leaves the realist with a dilemma: either accept a perspectival realism, or embrace skepticism. We shall explore this point in more detail in later chapters.

In more restricted forms, the argument from failure of fit has cropped up frequently throughout the history of philosophy: generally, the response is to locate some properties to which it does not apply, take a realist stance towards those, and then treat the remainder as either illusory or as reducible to the real properties. The

primary/secondary quality distinction is perhaps the most notable example of this move. Perception represents colour as an intrinsic property of objects. On analysis, it turns out that this view is unsustainable. The language of colour is not the language of the world. Having specified a particular way of carving up colour space, we can say that certain statements are true of false – so it is true, from a particular visual perspective, that grass is green – but there is no particular reason why any given visual perspective should be favoured over any other. The colours distinguished by the human’s visual system, under such-and-such conditions, are not more or less correct than the colours distinguished by the dog’s visual system, under such-and-such conditions. In this case, however, we take it that we can “escape from the colour perspective” and describe the situation in terms of surface reflectances, light, and the operation of a particular visual system. Certain types of surfaces, under a certain range of conditions, have the power to produce certain types of experiences in human subjects. The perspectivist challenge is that failure of fit generalizes to everything.

The failure of fit argument will occupy us for much of the rest of this thesis. The chapters are structured as follows. In the remainder of Chapter 5, I will elaborate on the failure of fit argument. In Chapter 6, I turn to realist responses to the argument. The perspectivist argues that underlying traditional realism is the implausible idea that the world is “lexicon-dependent”. An objector might deny that realism has ever been committed to such a thesis, or she might try to show that lexicon-dependency is more plausible than the perspectivist supposes, and that we can show that our representations to “fit” the world in the relevant sense. Several responses along these lines will be outlined in Chapter 6. However, by far the most sophisticated response is the appeal to natural kinds. It has long been supposed that there is some sense in which our best theories “carve nature at its joints”. This response will be discussed in Chapter 7. Next, in Chapter 8, I will outline how a perspectivist should understand taxonomies and kinds, once we give up the commitment to a language of the world.

5.2. Taking a perspective on Putnam's world

I claimed in the previous section that the failure of fit argument is *prima facie* more promising than the incompatible models argument, insofar as it seems to avoid falling into either straightforward realism or antirealism. But is this right? On one interpretation of the failure of fit argument, the problem arises just as forcefully, so it is important to clarify exactly what the failure of fit argument involves. In particular, it is worth distinguishing two approaches to perspectivism that are often conflated.

A natural interpretation of failure of fit is that failure of fit is a matter of *distortion*. That is, our theories never “mirror” nature, but always provide a distorted representation of nature, where the content of our theories is always conditioned by the specific features of our cultural and historical context. We have already seen this notion of distortion at work in the perspectivist interpretation of instruments. But just as in the case of instruments, and just as in the case of the incompatible models argument, to interpret failure of fit in terms of distortion raises the spectre of Chakravartty's dilemma. It seems that we can only be justified in talking of some model M providing a distorted representation of system S if we have some further representation of S that is independent of the model. Either we are able to specify the nature of the distortion, which allows application of Escape from Perspective, or we are not, in which case, we seem to be committed to antirealism about models of S. In the latter case, we seem to be claiming: S is not actually how M represents it to be, but we don't know what it really is like.

In talking of distortion, we are implicitly assuming a perspective-independent view. We are assuming a viewpoint from which the relation between the model and the target system can be evaluated, and the differences seen. This is an acceptable move in specific cases, since we often need to note ways in which particular models depart from the facts. But as a general philosophy of science, it is hard to see how this furnishes the resources for an alternative to standard realist/antirealist positions. If we know that we are distorting, and we can specify the nature of the idealization involved, then we must know something about the non-distorted truth: if we can specify that we are considering an idealized model without friction, then the target system is not frictionless. Notice, however, that the point about distortion is

tangential to the relativist approach developed in the previous chapter. On this approach, what is central to perspectivism is the notion of truth-from-a-perspective. To say that we have access only to perspectival facts or perspectival truths, that the true propositions uncovered by the sciences are true only within or relative to a perspective, need not involve any commitment to the notion that even our best theories and models necessarily distort.

Indeed, there is a tension between these two approaches. Kuhn says that evaluation of truth-values occurs only with a lexicon already in place, and he casts doubt on the notion that lexicons could be built into the world itself. Now, what could it mean for M to give us a distorted representation of S? We must be assuming that there are relevant facts about S that are only approximately captured, or are somehow transformed, in the representation given by M. This is understandable if, to put it in Kuhn's terminology, we think of the M-lexicon as deviating in certain ways from the S-lexicon. In any case, however we frame the matter, there must be nonperspectival facts or nonperspectival truths that are never quite adequately captured by our models. By contrast, if we claim that certain propositions drawn from our model M are true relative to perspective P, then from the point of view of P, there is no distortion of S. M simply provides an accurate representation of S. If we are concerned about distortion, we might be driven to talk of approximate truth instead of plain old truth. But it would remain the case that the propositions are approximately true relative to perspective P, and that there is no distortion to the extent that they are true. The distortion enters with the aspects that are judged to be false.

The distortion argument sees the relation between models and the world as perhaps like the relation between a photograph and a rushed painting of the photograph. The painting will inevitably differ in almost every way: the colours will be slightly different, the shapes less sharp, certain features will not be preserved, and so on. But we can still identify gross similarities, and we can treat the painting as a distorted representation of the photograph. They both involve the same "language" of shapes, lines, colours; with failure of fit arising from the fact that none of the "characters" in one exactly match any of the "characters" in the other. As I see it, the Kuhnian failure of fit argument proposes a more radical failure of correspondence. Kuhn's point – and this seems to be Giere's point at times also – is that there is no "language of the

world” for our models to match or fail to match. It is not so clear how Escape from Perspective can get any purchase here (I will consider this in more detail later; all I will note for now is that there is some prima facie reason for thinking that failure of fit has better prospects for supporting perspectivism than incompatible models).

An early instance of the failure of fit argument is provided by Putnam, in his example of the world of three atoms (Putnam 1996: 24). Putnam imagines a universe containing nothing but three simple atoms, call them “x1”, “x2”, and “x3”. He then asks: how many objects does this universe contain? The answer to this depends on the principles of material composition we use when counting the objects. Here are two principles. According to *unrestricted composition*, any two objects compose a further object, no matter what the properties of the objects are, no matter how they are related to each other. So, there is an object composed of my computer plus Frank Zappa’s left hand plus the brightest star in the Andromeda galaxy. Given this principle of object-individuation, Putnam’s world contains seven objects:

UC: x1, x2, x3, x1+x2, x1+x3, x2+x3, x1+x2+x3

At the other extreme, according to *compositional nihilism*, there are no circumstances under which two objects compose a further object. Putnam’s world contains three objects:

CN: x1, x2, x3

Finally, various forms of *compositional restrictivism* hold that objects sometimes, but not always, compose a further object. Perhaps, for instance, two objects compose a further object only when they are in physical contact. In this case, the number of objects in Putnam’s world depends on how exactly x1, x2, and x3 are related.

This brings us to debates concerning the metaphysics of material composition, which is somewhat beside the point the here. Whatever we think the composition relation is *really*, we do have alternative frameworks for carving up the world, and we may have very good reason to apply different frameworks in different contexts. Suppose we accept compositional nihilism, so that we take it to be a fact of the matter that there

are no objects with parts. Even so, we still need to say something about our ordinary use of the term “object” and terms which seem to denote objects with parts. Instead of truth-conditions, the nihilist will appeal to assertibility-conditions. They will say something along the lines of this: ““There is one chair in the room” is not true, but assertible just in case there are particles-arranged-chairwise located inside particles-arranged-roomwise.” In this way, we construct useful kinds.

In Putnam’s view, there really is no fact of the matter one way or the other. The world supports many alternative conceptualizations; and the world cannot be thought of or described independently of a conceptualization. Applying this to the failure of fit argument: the idea of perfect fit is problematic, because we have on the one hand a model, and on the other hand... what? What are we comparing the model to? Well, the world. But to make any comparison, the world must already be conceptualized.

It is tempting to suppose that there is simply nothing interesting to be said about Putnam’s world, or principles of material composition more generally. We might follow Stanford (2017) and view questions about material composition as merely verbal disputes, not substantive disagreements. Once we have specified the three atoms in Putnam’s world, and described their properties and relations, we have said everything that there is to say. There is simply no need to be concerned about how to divide the world up into “objects” per some technical definition of the term. We can define “object” however we wish here, and the definition that we should adopt is simply a matter of pragmatic factors. The question is simply: which provides us with the smoothest and most convenient conceptualization of the world? The principles of material composition do not reveal different “perspectival facts”: they reveal precisely the same facts, using slightly different terminology. But as Putnam points out, this move looks question-begging. Why should we accept that, in stating that (for example) Putnam’s world contains three atoms, we have said everything about it that there is to say? The question here is what exactly constitutes a “complete inventory” of Putnam’s world. Now obviously, if we are considering the world through the lens of unrestricted composition, the inventory listing just x_1 , x_2 , and x_3 does not say everything there is to say about it.

An interesting approach to this problem has been proposed by Hirsch (2005). According to Hirsch, the different conceptualizations of Putnam's world can be seen to be merely verbal disputes when we consider the truth conditions for the different statements. UC tells us that there are seven objects; CN tells us that there three objects. We can take each statement as describing a class of possible worlds. That is, for each statement, there is some class of possible worlds in which that statement is true. Then "sentences have the same truth conditions if (relative to any context of utterance) they hold true in the same possible worlds" (Hirsch 2005: 72). It follows that in the debate about Putnam's world, the two sentences have the same truth conditions. UC's "there are seven objects" and CN's "there are three objects" are true of the same possible worlds. This is perhaps not a surprising result: when introducing the debate between UC and CN, we were asked to imagine Putnam's world, not Putnam's worlds. They are two ways of carving up the same world. From a third-person point of view, we take their statements as true under the same conditions; we must do this to set up the problem in the first place.

Again, however, this move looks question-begging. The realist claims that "there are seven objects" and "there are three objects" are true of the same world. But why should we count these as *the same* world? More generally, how exactly are we individuating possible worlds here? From the point of view of unrestricted composition, the compositional nihilist has not specified the same world, since the compositional nihilist is supposing that there three objects and that those objects are not parts of further objects. Of course, from the point of view of unrestricted composition, the compositional nihilist has not specified any *possible* world at all – at least, not any metaphysically possible world, perhaps not even a logically possible world. Perspectives shape not just what exists, but also the space of possibilities. It should not be surprising that, according a perspectivist, the modal facts are also perspectival facts.

So far, then, a draw. From the perspective of UC, it is a fact that there are seven objects. Is this, as a perspectivist would have it, a perspectival fact, inaccessible from CN, or is the traditional realist right that we simply have different ways of describing the same perspective-independent fact? I'm not sure that the answer particularly matters, since either way, I don't think Putnam's world sustains any

interesting kind of perspectivism. The primary reason is simply that, even if we take it that object-frameworks generate perspectival facts about the objects in Putnam's world, there is no reason to treat the initial description, that the world contains three atoms, as perspectival. What we have in Putnam's world is a perspectival account of object-frameworks specifically. We can explain the truth of statements such as "there are three objects", under CN, in terms of (a) a non-perspectival description of the atoms contained in Putnam's world and (b) the characteristics of the CN perspective. A traditional realist would say that the lexicon of Putnam's world is a lexicon of atoms, and the truths about those atoms ground the truths of claims made in terms of object-frameworks. This is analogous the application of a primary/secondary quality distinction to deal with colour.

Putnam's world does, however, illustrate two important points. First, the problem that motivates Putnam's thought experiment is a kind of mismatch between our representational tools and the world. Within any given object-framework, we can make true claims, yet the world itself does not determine how we carve out its objects. Second, alternative object-frameworks do not in any obvious way introduce distortion or idealization. From a perspectivist point of view, both UC and CN provide access to facts that are relative to those frameworks, but neither UC nor CN involve distorting the world. Nothing false is said about the world via either framework. It is simply that if we want to talk about objects and individuate objects, we need some framework for doing this. Recall Chakravartty (2010): "A perspectival fact is a proposition that is only true from, or within, or relative to, a given perspective." The proposition that Putnam's world contains seven objects is true relative to UC, so that Putnam's world contains seven objects is a perspectival fact about it. The truth is constituted within the perspective.

5.3. Complexity

Putnam's world is simple. By contrast, what grounds the contemporary perspectivist argument is the notion that the world is extremely complex, and that it far exceeds our cognitive powers to represent any worldly system in its totality. It is easy to see how this leads to treating perspectivism as a matter of distortion. Since the world is complex, we are forced, when constructing models of worldly systems, to simplify.

Our representations are inevitably partial, abstract, vague, and idealized in various respects. As we have seen, this kind of argument is explicit in both Giere and Teller. Having distinguished the relativist and distortion approaches, we will now develop this point about complexity in a somewhat different direction.

Putnam's world provides an initial illustration of the relativist approach to perspectivism, in contrast to the distortion argument – and this of course dovetails nicely with the arguments of the previous chapter. In Putnam's world, if we accept the perspectivist analysis, then perspectival facts arise from a kind of underdetermination. The world does not determine the facts about object-parthood relations; these facts arise only once we conceptualize the world in some way. The problem with Putnam's world as an argument for perspectivism, beyond the fact that there is no reason to take the specification of three atoms as perspectival, is that we can resolve the tension in different conceptualizations by introducing a disjunction. For Putnam's world, we can state a disjunction. Either UC, on which there are seven objects, or CN, on which here are three objects, or... (add your favourite theory of object-parthood relations here). Either description works just as well as the other in terms of accommodating the non-perspectival facts, so if we are simply concerned about the accuracy of our representations, then both are acceptable and there is no further debate to be had. Both provide equally accurate pictures of the world, but we might evaluate them differently per certain pragmatic standards. Here are two options for individuating objects: take your pick. In the disjunction, we have a straightforwardly correct description of all the relevant facts about Putnam's world. But this move is not available to us. We cannot step back and provide a disjunction of possible frameworks in the same kind of way.

Why not? Notice that when thinking about Putnam's world, we hold two things fixed. First, we hold fixed certain facts. We are assuming that the world contains three atoms. We are entitled to make this assumption, since Putnam's world is just a thought experiment. But what about our own world? In this case, we are not entitled to make any assumptions such as this – not in general, across all contexts of inquiry, at any rate. There are various contexts in which we grant certain assumptions: as Kuhn would point out, progress in normal science often requires scientists to put aside skepticism about prevailing theories and just work on developing the theories.

Outside of a scientific context, all of us hold various beliefs that we recognize are open to skeptical attack. So we do, perhaps, make assumptions that are not fully justified. But the appropriate assumptions depend on what activities we are participating in and what kind of questions we are asking. This brings us to the second factor we hold fixed in thinking about Putnam's world. Putnam introduces this thought experiment to ask us the question: How many objects are there in this world? This is the only question we are interested in, the only reason why Putnam introduced the thought experiment in the first place. Our only standard by which to evaluate different object-frameworks is in terms of how well they answer this specific question.

When inquiring about the real world, there is not in general any fixed picture of things that we apply different conceptualizations to; there are simply a host of different conceptualizations. This point will be substantiated in later chapters, but an intuitive presentation of it can be made by returning to the perspectivist's colour analogy. How many colours are there? There are countless ways of carving up colour space. Different languages draw the lines in different places and group colours differently. Even within English, there are various standard ways of dividing up colour space, depending on how fine-grained we want to carve it. Now, given a particular conceptualization of colour space, we can understand alternative frameworks in terms of our own framework. For example: language L1 uses a single term "a" for what we call both "blue" and "green"; language L2 uses a single term "b" for the warm colours "white", "yellow", "red", etc. and single term "c" for cool colours "blue", "green", "black", etc. In this way, a competent user of one colour-framework can come learn other frameworks – just as how a competent user of one object-framework can learn to translate their statements into another one. In the case of colour, however, there is no neutral description. Putnam supplies us with the description that the world contains three atoms. Nothing supplies us with any description of the colours of the world. Of course, we all perceive the same world, but perceiving is not describing.

When inquiring about the real world, there is not just one fixed question that organizes our research. This point is of great significance. If we take the line of relativistic perspectivism that truth is constituted within a perspective, then we cannot

hold truth to be the general aim of any perspective. That is, truth is not a perspective-independent value against which various possible perspectives are judged.

Achieving the truth is not the criterion of success of a perspective. To put it in Elgin's terms: what makes for a good perspective is determined by various more-or-less pragmatic factors, and then within this perspective, we can state various truths.

However, the factors that make for a good perspective are not eternal. In Putnam's world, we take it that we are answering one question. In the real world, the questions that drive inquiry are constantly in flux. This is partly a consequence of complexity.

As Wray argues, because our cognitive powers are so limited relative to the enormous complexity of the world, scientists are forced to ignore certain features of target system and so they can construct only partial representations of those systems (Wray 2018, Chapter 12). Exactly which representations would be appropriate depends, of course, on what our goals are. How good of a job a theory does, the degree to which it exhibits various theoretical virtues, is dependent on our research interests, and these are subject to change. Wray gives the example that early modern astronomers attempted to account for the motion of the planets and for the relative brightness of the planets at different times. They did not attempt to account for the colour or mass of the planets. These features were disregarded. Of course, there were good reasons for this: for example, they simply had no reliable way to work out the mass. With the rise of Newtonian mechanics came the possibility of estimating mass, and the attempt to unify terrestrial and celestial phenomena under a single theory; in this context, mass became a relevant feature that planetary models were expected to account for (Wray 2018: 190). Along the same lines, Wray notes that since Aristotelians took the planets to be embedded in celestial spheres, they had no trouble explaining the stability of planetary orbits. When the Aristotelian cosmology was displaced, the stability of planetary orbits became a pressing research problem (2018: 201). The question, "why do the planetary orbits remain stable?" was not a topic of investigation for astronomers in the Aristotelian tradition. It took a great deal of conceptual and practical development in astronomy before this was seen as an important question.

It is often the case in the history of science that phenomena that are initially regarded as intriguing though minor anomalies later become the lynchpins around which new

theories are constructed. Initial attempts to explain the zebra pattern of stripes of magnetization on the seafloor took specifically the *positive* magnetizations as the explanandum; the fact that some areas were negatively magnetized was not seen as significant. Yet the alternation between positive and negative magnetization became central to plate tectonic theory when it was explained by combining the hypotheses of geomagnetic reversal and seafloor spreading (Schindler 2011). Negative magnetizations had long been known to exist, but models were not required to account for them.

Obviously, theories and models are expected to accommodate the phenomena. But what this amounts to in general is that the phenomena should not be known to contradict our theories. If our geophysical model predicts that there should only be positive magnetizations, and no negative magnetizations, we have an anomaly, which there would be at least some pressure to resolve. It is another matter entirely to construct explanatory models of a given phenomenon. Various phenomena will be left unexplained or will be seen as fitting smoothly into a broader theoretical scheme in such a way that no detailed explanation is called for. Geophysicists in the 50s and early 60s may have assumed that the negative magnetizations could be caused by factors such as seismic activity along fault lines (Schindler 2011: 47). The origin of stripes of magnetization was a puzzle, but given the existence of positive magnetization, the alternation of positive and negative magnetizations could straightforwardly result from casual mechanisms that were thought to be well-understood.

Of course, research questions are guided by the background theories, and so in the cases discussed above, they were guided by what we would consider radically mistaken presuppositions. It is unsurprising that research questions change over scientific revolutions. However, the wide scope of potential research questions arises not just from background theories, but from the complexity of world, which forces scientists to narrow their focus to specific phenomena. Alternative research questions, producing alternative adequate representations, proliferate within a single research tradition. Consider the boiling point of water. What is the boiling point of water? Assume for the sake of argument that we already have reliable thermometers, but we have not yet examined water (so we are putting aside the

many complexities involved in constructing a stable and reliable practice of thermometry; cf. Chang 2004). Given this assumption, the obvious way to answer our question is to take a series of thermometer readings, plot them on a graph, and then find a line of best fit. The readings will show a scatter around 100°C.

This will lead to new research questions, such as: Why does water boil at 100°C? Or: Why does the boiling point of water differ from the boiling point of mercury? The fact to be explained here is the invariant boiling point of water. However, there are a variety of other, more specific phenomena that scientists might choose to investigate, or that might become relevant to the evaluation of a theory (McAllister 1997). The boiling point of water is dependent on a host of factors: atmospheric pressure, impurities in the water, the proportion of water molecules composed of deuterium, the structure of the container, the materials of the container, the interaction with the thermometer, and so on. Changing all of these factors will result in different boiling points, and of these might be taken as the explanandum. The concept “the boiling point of water” is an idealization: to get an invariant boiling point, we would need a perfectly pure sample of water, in precisely fixed environmental conditions. This never obtains in the real world. Of course, if we are interested in explaining the difference between the boiling point of water and the boiling point of mercury, it would be natural to at least initially take the idealization of an invariant boiling point as the relevant phenomenon. But there are a variety of other questions we can ask, e.g.: how does the boiling point of water change with increasing salinity?

Persectivists often raise points along these lines to elaborate the ways in which our models inevitably distort (e.g. Teller 2019). As I have emphasized, distortion is not really the point here. I think that we know various truths about the boiling point of water, and I am not sure how to square this with the idea that all propositions involve some distortion, since “distortion” just seems to mean “false in some respect”. My claim is that the complexity of the world has two consequences: first, it forces us to construct only partial representations of any given worldly system, and second, it leads to an abundance of properties that we could take as explananda. Whether or not a representation is appropriate depends partly on the specific questions we are asking about a given system. In different contexts, a variety of different representations will be appropriate.

In Putnam's world, with its fixed picture of three atoms, and fixed question guiding our inquiry, of how many objects the world contains, we can resolve the problem of alternative conceptualizations by offering a disjunction. We can say: here is UC, and here is CN; we can easily translate from one to the other; and they are both equally adequate in the following sense: they both accommodate the fixed set of facts, and they both provide a simple, intuitive, and easily applicable answer to the question that we are trying to ask. There may be no way to decide between them, but there is also no pressure to decide. Even if these alternative frameworks do generate perspectival facts about object-parthood relations, once we state the disjunction, we have said all that there is to say.

For our inquiries into the real world, by contrast, we cannot make this disjunctive move. There is no fixed background picture and no fixed question guiding inquiry. The disjunction could not simply state two equally adequate theories. It would also need to state the broader context motivating those theories. That is, we cannot simply say: T1 or T2 or T3 or... Rather we have: C1/T1 or C2/T2 or C3/T3..., where "C1", "C2", "C3", etc., specify various background assumptions, standards, and research questions. This disjunction cannot be completed, because we have no way to know what kind of factors will become relevant to scientists in the future. As we have seen, changes to research questions are often a product of theoretical revolutions, or of new discoveries, or of increasing precision in the use of instruments. Any new discovery or new methodological development is liable to prompt scientists to focus on new phenomena. Since we have not yet conceived of such developments, we cannot complete the disjunction. Furthermore, even if we could complete the disjunction in principle, we could never actually occupy all of the different contexts. Today we can understand the concerns that motivated Aristotelian astronomers, but little of this is relevant to contemporary astronomical research. For one thing, we cannot simply ignore the discoveries of the last few hundred years. The Aristotelian perspective is not in practice open to us as a framework for organizing inquiry.

According to the failure of fit argument, there is a kind of mismatch between our representational tools and the world, which entails that truth is constituted within a

given perspective. Putnam's world is a useful illustration of how this argument works. Unfortunately, it is vulnerable to the objection that the conflict between competing object-frameworks is nothing more than a verbal dispute. But even if scientific inquiry similarly produces frameworks, with truth constituted within those frameworks, there are good reasons to think that the appeal to verbal dispute fails to generalize.

It is important to make clear that the argument here is not that truth is simply determined by whatever is proposed in any theoretical perspective that some people consider useful in some context. For example, I am not claiming that it is a perspectival fact that the planets are embedded in crystal spheres, that it is true-from-the-Aristotelian-perspective that the planets are embedded in crystal spheres. In general, there are often good reasons to suppose that other people are simply mistaken about various things, even from the point of view of the theoretical perspectives they work in. All that follows from the failure of fit argument is that the lexicons that are used to construct true (and false) propositions are supplied by the theoretical perspective, not the target systems in the world.

5.4. Incompatible models again

Before concluding this chapter, it is worth returning to the incompatible models argument. I have suggested that the argument from incompatible models does not have adequate resources to motivate perspectivism. But it might be objected at this point: is not the failure of fit argument just another kind of incompatibility argument? Rather than talking about models, we are talking more broadly in terms of languages and lexicons. While many perspectivists have appealed to the way in which scientists generate multiple incompatible models of worldly systems, we are now appealing to the multiple incompatible representations, more generally speaking. Indeed, it might be the case that the incompatibility of different models partly consists in their use of different lexicons.

There are, however, significant differences between the incompatible models and failure of fit arguments. Most importantly, as I noted earlier, we cannot step outside of our lexicons in the way that we can do so for incompatible models. Incompatible models are straightforwardly accommodated in other positions, and so this argument

fails to provide a grounds for perspectivism. There is prima facie reason to think that failure of fit has better prospects here, and it will be the task of the next chapters to substantiate this impression. Moreover, the issue here is not incompatibility. Even if in some domain there is only one theory that we take to be uniquely the best, only one theory that strikes a good balance on the theoretical virtues, we can argue that the truth cannot be perspective-independent. Indeed, in practice, scientists do often converge on one specific theory or model for a given domain or target system. The alternative frameworks of Putnam's world are merely illustrative. In practice, developing even just one adequate framework can be extraordinarily difficult. Consider again the colour analogy. It is logically possible that, given the structure of our visual systems, and given a particular way of life, there is exactly one uniquely best way of carving up colour space. Yet it would remain the case we do the carving, and that the truth of a statement such as "grass is green" is constituted within the colour-perspective in question. Again, this point will be substantiated further in later chapters.

Although I have been critical of the incompatible models argument, I do want to point out that incompatible models may still have a role to play in the justification of perspectivism. This is because we might appeal to incompatible models to strengthen the failure of fit argument. In particular, the perspectivist could appeal to incompatible models as part of an inference to the best explanation for failure of fit. Here is a sketch of that argument, altering Cretu's presentation of the original incompatible models argument:

(P1*) Diverse modelling practices are ubiquitous in successful scientific practice.

(P2) Diverse modelling practices often give rise to multiple, different models of the same target system.

(P3*) The best explanation for (P1) and (P2) is that models use constructed lexicons.

(P4) Knowledge obtained through constructed lexicons regarding the same target system is inherently perspectival.

(C) Therefore, modelling practices yield perspectival knowledge about the world.

In any case where we have incompatible models of some phenomenon, we can account for this in standard realist or antirealist terms. But then when considering the ubiquity of incompatible models, or the use to which they are put in scientific practice, it might be held that a better explanation is furnished by taking models to involve lexicons that are constructed, rather than lexicons that mirror the world's lexicon. Here is one way to develop the argument: Realism, the perspectivist might hold, provides a poorer explanation of the success of incompatible models because the realist cannot provide any *general* explanation. The realist must consider each specific case of incompatible models, and will make different moves in each case: propose an underlying theory that resolves apparent incompatibility, or take one model to be epistemically superior to the others, or take the models to be representing different things... and so on. Antirealism, by contrast, provides no explanation of success at all; and it also runs into the problem that we sometimes have incompatible representations of observable phenomena, and so threatens to fall to radical skepticism. Perhaps, then, the perspectivist explanation has better prospects. However, there would be no reason to take the perspectivist explanation seriously, in the absence of the more general failure of fit argument. So, I will not explore this version of the incompatible models argument any further here.

In this chapter, I have outlined an argument for perspectivism, the argument from failure of fit. I believe that this argument has better prospects for supporting perspectivism than either instruments or incompatible models, and it will occupy us for the remainder of this thesis. In the next chapter, I consider some initial realist responses to the challenge of failure of fit, and then Chapter 7 takes up one of these responses in detail, the appeal to natural kinds.

Chapter 6

The Language of the World

6.1. Introduction

The perspectivist challenges the realist to show that there is the right kind of relation between our representations and the world. Could there be a language of the world, and could our theories speak it? While Kuhn and Giere find the notion of a “lexicon-dependent” world vague or perhaps incoherent, there are reasons to take this idea seriously. In this section, I will consider how traditional realists might initially defend it, before turning to the perspectivist attack. There are broadly speaking two ways for a realist to approach the failure of fit challenge. First, the realist might give general reasons for optimism that our representations match the world. Second, the realist can try to specify exactly what the “language of the world” is.

I will begin by considering reasons for optimism. The only extensive discussions of this that I am aware of in the literature take the form of indirect arguments: the realist can proceed by arguing for some particular version of metaphysical realism; if the considerations favouring this position significantly outweigh that of its rivals, then perhaps we can reasonably view the failure of fit challenge as an anomaly that can be put aside, even if we do not have a direct response to it. That is, we can grant that our position has unsolved problems, but still be optimistic that those problems can be solved, if there are strong enough arguments in favour of our position. To this point, I can only say that one of the central goals of this thesis is precisely to undermine traditional realism, so I do not find this response convincing. With that said, I will consider three arguments for optimism that directly respond to the challenge. As noted, however, this remains an under-explored topic in the literature.

6.2. Denial

One option for the realist is simply to deny that there is any challenge here in the first place. All that has ever been required for realism, it may be said, is that we are justified in believing that at least some of our best theories are true. The realist does not need to make any further commitments: in particular, the realist does not need to

adopt any controversial positions with respect to metaphysics or theories of truth. This kind of response is similar to that of Musgrave (1989), responding to Fine's natural ontological attitude. Fine (1986) attempted to carve out a *via media* between realism and antirealism, arguing that we can identify a "core position" shared by realists and antirealists. This core position is simply that well-confirmed theoretical claims are true, on par with the "homely truths" about cars and mountains and oceans and various other things that we perceive with the senses. All parties in the debate, Fine claims, accept that there are, for example, electrons, and that contemporary physical theory is what justifies this. Just as we all trust the evidence of our senses, we all trust the results of science. What distinguishes realists and antirealists is that they each add controversial philosophical claims to this core position: the realist may adopt a correspondence theory of truth, while the antirealist may adopt a verificationist theory of truth. Fine's natural ontological attitude, by contrast, simply accepts the core position, and nothing more. We do not need to enter any metaphysical debates; we can simply trust the results of science.

Musgrave objects that the debate between realists and antirealists is precisely a debate about whether we are justified in taking our best theories to be true. Antirealists do not treat "homely truths" about observables, and theoretical claims about unobservables, as being on a par, so they simply reject what Fine calls the "core position". Indeed, the core position is, Musgrave argues, all you need to be a realist. Similarly, I doubt that Musgrave would be moved by the perspectivist argument outlined earlier. Realists are not committed to the view that truth consists in "correspondence" between theories and the world, nor that there is a "language of the world", or anything along those lines. The perspectivist is responding to a position much stronger than what any realist ever needed to accept.

Musgrave is half-right. There are plenty of antirealists who deny that we are justified in taking claims about unobservables to be true, and who deny that observables and unobservables are on a par. Such brands of antirealism are probably the most influential in contemporary philosophy of science. However, Musgrave is overlooking other forms of antirealism, particularly those in the social constructivist tradition (see Turner 2007: 130-162). Social constructivists are generally quite happy to grant that various scientific claims are true, that we are justified in believing them to be true,

that we know them to be true, even that they are true of the external world. Many social constructivists will also treat everyday “homely truths” and scientific truths as on a par, viewing both as products of social construction. What distinguishes social constructivism from realism is that they each add a different account of truth to the core position. An example from Turner (2007: 147) makes this clear. Suppose we find a set of tracks in the snow, and we propose the hypothesis:

H: A deer passed this way not long ago.

We then accumulate a great deal of evidence for H. Perhaps we find lots of deer droppings nearby, then we discover that the area has a high deer population, then closer inspection of the tracks reveals patterns only known to be caused by deer, and so on. Both realists and constructivists will accept that H is true. This is part of the core position. It also seems that H is a “homely truth”, so we can put aside questions about the justification of theoretical posits beyond the observable. What realists add to H will be something along these lines:

HR: A deer passed this way not long ago, and the deer’s making of the tracks in the snow occurred independently of the mental.

For the realist, the fact that the deer passed this way not long ago is a fact about the world, that is discovered through our investigation. The constructivist addition:

HC: A deer passed this way not long ago, and the deer’s making of the tracks in the snow was dependent, in one way or another, on our thoughts now.

HC is a striking claim and initially may seem hopelessly implausible. Some caution is required in interpreting it. Most importantly, HC should not be read as making a claim about causal connections. The claim is not that our thoughts caused the deer and its tracks in the snow to come into existence. There is no compelling theory of the world on which that is true. In any case, from a constructivist point of view, we also construct the causal connections – that is, the causal connections between things are also dependent, in one way or another, on our thoughts now – and clearly, we have not constructed things such that our thoughts cause the existence of deer. So

how should we think about HC? We might draw on something like the failure of fit argument. Here is a crude way to defend something like HC: Facts are true propositions, or correspond to true propositions, or something along those lines. A proposition is stated in terms of a lexicon. But the lexicon is constructed by us. We determine what counts as a deer, what counts as snow, what counts as tracks, etc. As Woolgar (1988) puts it: “there is no sense in which we can claim the phenomenon ... has an existence independent of its means of expression ... The organization of discourse is the object. Facts and objects in the world are inescapably textual constructions.” (Quoted in Kukla 2000: 18.) HC therefore makes a constitutive, rather than a causal, claim. The facts about the deer are constituted, or at least partially constituted, by our thoughts now. This includes the fact that the behaviour of the deer is causally independent of our thoughts. That is, if we were to model the behaviour of the deer, then depending on what exactly we were interested in explaining, we would appeal to features of its environment, such as facts about the forest in which it lives, features of its body, features of its evolutionary history, and so on, not the features of human minds.

Even with this clarification, we might quibble over the characterization of social constructivism. The constructivist may object that she is not committed to supposing that events in the world are dependent on our thoughts in any sense of “dependent”. But this is only illustrative; constructivists are committed to something *like* HC. In any case, the point is that if realists are only committed to the core position, and nothing more, then they have no way of ruling out HC and similar claims. This would be a weak form of realism – surely weak enough to include perspectival realism. In order to understand the debate between realists, constructivists, idealists, perspectivists, and so on, we need to go beyond the core position.

What does traditional realism involve? At the very least, realists assert not simply that our best theories are true, but that our best theories are true of a mind-independent, external world. Furthermore, truth itself must in some sense be independent of us. It is true that there are electrons, and this would remain true, even if nobody believed that there were electrons, or even if nobody ever existed in the first place. Truth is transcendent in that there may be truths such that it is, in principle, impossible for us to know them. This point is important: what makes

propositions true is not our available methods of verification (as a verificationist might say), or processes of social negotiation (as a social constructivist might say), but how those propositions relate to the mind-independent world. For H to be true, it must be the case that “deer” picks out some object in the world; and while, obviously, we fix the meaning of the term “deer”, the realist will insist that given a particular meaning, we have no control over whether this term has a referent, nor what precisely the referent is. I don’t think that this is controversial as a characterization of the commitments of traditional realism. But with this in place, the perspectivist has enough to forward the failure of fit argument.

Denial is not a response to failure of fit; it is simply a refusal to engage with it. There is nothing necessarily irrational about this: we can always invoke Fine’s natural ontological attitude, point out that we all share the core position that various theoretical claims are true, and then refrain from elaborating on what exactly truth consists in. If we are arguing against those antirealists who do not share the core position, such as van Fraassen, then the distinctions between those who do may be of less importance. But there *is* a distinction between traditional realism and perspectival realism, and the failure of fit argument is relevant to it.

6.3. Approximation

Realists do not claim that our best theories are true, but only that they are approximately true. Can this provide a response to the failure of fit challenge? The thought might be something like this: since realism only requires a commitment to approximate truth, realists need not hold that there is any kind of match between the lexicon of our theories and the world itself. Perhaps a true theory must be a theory that gives “a description in the object’s own terms or language,” but approximately true theories face no such constraint.

Although this line has some initial intuitive appeal, further consideration shows it to be implausible. What can it mean to say that a statement is approximately true? One simple option is to suppose that we have a conjunctive statement, where only one part of the conjunction is false: A and B and C and..., where only A is false. Perhaps this is a long conjunction of propositions describing a car, and A is “the car is red”,

when in fact the car is green. Our conjunction represents the car as being a different colour than it is. It is approximately true, because it is true in all other respects. Now clearly, approximate truth in this sense does nothing to help the realist, because it is parasitic on straightforward truth. The conjunction inherits its approximate truth from the truth of the propositions other than A.

Alternatively, we can think of the approximate truth of a proposition in terms of its closeness to the truth. “London and New York are 3500 miles apart” is approximately true: the more accurate distance is that they are 3459 miles apart. Of course, even this more accurate estimate is not exactly true. Or take, “the moon orbits the Earth”. This is technically false because the moon orbits the barycentre of the Earth-moon system, but it approximately describes the behaviour of the moon, in the sense that an orbit centred perfectly on the Earth is very close to the moon’s actual orbit. Approximate truth in this sense provides no comfort for the realist. Again, approximate truth turns out to be parasitic on straightforward truth. What makes it approximately true that London and New York are 3500 miles apart is the truth that London and New York are 3459 miles apart... or something like that. As noted, it is not true that London and New York are 3459 miles apart. It is not at all obvious what the truth is that is being approximated by these approximately true propositions. This, of course, is grist for the perspectivist’s mill: Teller (2018b) in particular has made much of examples such as this.

6.4. Evolution

A second reason for optimism is what I am calling the argument from evolution. For Kuhn and Giere, there is a mystery at the heart of traditional realism. The theoretical lexicon must correspond to the world. Clearly, languages and lexicons are constructed by humans. How could it be the case that a language or lexicon is built into the world itself? But to this, the realist might argue that given that humans evolved within the world, and the capacity for mental representation evolved, presumably, to represent the world, there is no mystery at all about why there would be a correspondence here. There is good reason to expect an appropriate match between our representations and the world, because our conceptual capacities evolved to match it. Consider this in terms of the theological view. God spoke the

world into existence, and he also spoke our minds into existence; presumably, when he did this, he intended for us to be able to discover truths about the world. It is no surprise, then, that our language would match the language of the world. The capacity of our mental representations to appropriately match the world is a kind of adaptation, and on the theological view, is explained by God bringing about both sides of the correspondence. Today, of course, we explain adaptation by appealing to evolutionary theory rather than the action of a designer.

Here is another way to put the point. Kuhn and Giere begin by thinking about theories and language, and then wonder how the world could itself have relevant lexicon-like features. After all, languages are human tools. In projecting the properties of language onto the world, we seem to be engaging in a bizarre kind of anthropomorphism. But the realist can suggest that we have the kind of representational tools we do because of the way the world is. If the world had not been lexicon-dependent, then any organisms within it would have developed different representational tools. The idea here is that the very fact that our representational tools have a conceptual structure is what gives us reason to think that the world itself has a conceptual structure; had the world been different, so our representational tools would have been different.

This is analogous to a point made by Peter Lipton (2002) in his response to Cartwright's claim that the world is "dappled", that there are no universally true laws but that laws only hold in particular regions. One of Cartwright's arguments for this is that we have succeeded in constructing lawlike models only for a limited range of cases, and even where lawlike models can be constructed, these usually require including a host of idealizing assumptions. For example, we have no model which could allow us to predict the trajectory of a bill thrown out of a window and fluttering to the ground. Cartwright is moving from a claim about our representational tools to a claim about metaphysics. Lipton raises a general objection to these kinds of arguments. What Cartwright would need here is a kind of counterfactual tracking: she needs to show that if the world had been fully law-governed, then we would have been able to construct lawlike models across the board. The problem for Cartwright is that it's not clear how we could show that, since we can also account for our failure to construct such models by appealing to our cognitive limitations. If the world had

been fully law-governed, we still wouldn't have been able to construct a lawlike model of the trajectory of the bill. So, the fact that we don't have such a model does not support the dappled world hypothesis.

We can expect a similar exchange between the realist and the perspectivist. The perspectivist argues that the realist is committed to the notion that the world is lexicon-dependent. But this, the perspectivist worries, is not an idea that can be made coherent. The realist now wants to offer an argument for optimism that the world is lexicon-dependent: since humans evolved capacities to represent the world, we should expect relevant matches between our tools and the world. It seems then that the realist and the perspectivist disagree about the counterfactual: If the world had not been lexicon-dependent, then our representational tools would have been different. It is extremely difficult to know where to begin with evaluating a counterfactual like this. But this is no matter, from the realist's point of view. All she needs to do is defuse the perspectivist's challenge. So it is now up to the perspectivist to explain why the realist cannot suppose that, had the world not been lexicon-dependent, our representational tools would have been different.

Of course, this response does not tell us exactly what the "lexicon-dependence" of the world is supposed to amount to. However, as noted, the argument from evolution is merely an argument for optimism that lexicon-dependence can, in principle, be spelt out in a way that is coherent and that vindicates traditional realism.

One objection to this argument is that it's simply question-begging. It assumes that all humans have occupied the same world; that all of humanity evolved in the same environment in relevant respects. If our ancestors evolved in worlds with different lexicons, then there would be no reason to assume that the lexicon of our language matches the lexicon of our world. But if facts are perspectival, then the facts, and presumably also the lexicons, may have been different in the past, when people had different perspectives.

There are two points in response to this. First, arguably, the realist only needs coherence by her own lights here. Failure of fit is presented as an argument against realism. The basic claim is that we cannot make sense of the notion of

correspondence between representation and world that is required by realism. The argument from evolution is supposed to give the realist a license for optimism, so now the realist can say that if realism is correct, then we can understand how it is that the correspondence obtains. Second, perspectivists should accept that all humans evolved in relevantly similar environments, insofar as perspectivists accept contemporary evolutionary theory. Again, recall the point from the discussion of the denial response. Both realists and perspectivists can agree on which scientific theories are true. They differ in their analysis of what truth consists in.

However, there are more serious problems for the argument from evolution. Most notably, our best science appears to reveal that we have been endowed with representational capacities that systematically generate false representations of the world. It is understandable why this should be so: what matters for evolution is not true representation but useful representation, where utility is a matter of survival and reproduction. Of course, organisms must respond appropriately to the world, but appropriate responses need not be mediated by representation of any kind, let alone representation that corresponds to the world.

Consider, for example, folk physics (Smith and Casati 1994), our untrained, everyday understanding of physics, which involves significant simplification and distortion of physical phenomena. Folk physics includes an assumption of absolute simultaneity: a reasonable assumption, given that relativistic effects are negligible in everyday life. But folk physics also contains distortions about matters that are relatively easily observable. McCloskey (1983) presented subjects with various problems about motion. In one problem, they are told that a person has a ball attached to a string, and is spinning the ball around above their head, so that it follows a circular path. Eventually, the string breaks, and the ball flies off. Subjects were asked to describe the path of the ball after the break. Rather than depict the ball flying in a straight line (the correct answer), many of the subjects assumed that the ball would continue to follow a curved path, which would gradually straighten out, as though it had been imparted with some intrinsic “circular motion” that would gradually be lost. McCloskey suggests that folk physics assumes a kind of impetus theory of motion.

Or consider visual perception, which seems to present objects with intrinsic colour properties. As we have already seen, colour properties are the exemplary case for perspectivists. There are powerful reasons to think that colour properties do not inhere in objects intrinsically, but are, at best, a product of interaction between organism and environment. In a wide range of cases, it appears that there is selective pressure for representation that abstracts, distorts, and simplifies. In addition to this, as our theories of the world have become more sophisticated, and have revealed more of its facets, they have also become increasingly less intelligible (cf. Dear 2006). Theories dealing with quantum phenomena violate common sense to the degree that it is not clear that it is even possible for us to understand them adequately. If the world has a lexicon, there is reason to think that it is not our lexicon. That humans evolved in the world does not help here.

I know of no other general argument for optimism, so I will now turn to specific proposals for what the language of the world could be.

6.5. Propositions

There is, of course, no question that languages are constructed by humans. As Giere points out in his discussion of the failure of fit argument, one natural response to this is to suggest that all languages are constructed in order to express abstract meanings, such as propositions (2015: 189). In English we say, “snow is white”, in German we say, “schnee ist weiss”; but both express the same proposition, namely that *snow is white*. The relevant fit is then not really between natural languages and the world, but rather between propositions and the world. This is indeed how many philosophers had traditionally thought about truth, especially in the tradition of correspondence accounts of truth.

The problem with this response is that, as Giere notes, the same difficulties arise whether we are talking about natural languages or the abstract propositions that natural language sentences express. There may well be good reasons to introduce propositions – perhaps we need them to account for how two sentences in different languages can express the same meaning, or how two people can share the same belief, and so on. But the challenge that is raised by the failure of fit argument is that

realism seems to require that the world bears a kind of conceptual structure. Of course, propositions have a conceptual structure. Simply postulating the existence of some abstract objects that bear conceptual structure does nothing to explain how the world itself could have this structure. All we do by introducing propositions is shift the problem from language/world correspondence to proposition/world correspondence. To approach this point from a different direction, consider that nobody supposes that states of affairs in the world are literally just propositions, as propositions are understood in this context. Propositions are supposed to be abstract entities that are the bearers of truth-value (McGrath and Frank 2018). Propositions about a mountain might be true or false but, it would be an obvious category error to suppose that a mountain, just in itself, is true or false.

6.6. *Mathematization*

Perhaps the most famous instance of an appeal to the metaphor of the language of nature is Galileo's claim that the book of nature is written in the language of mathematics. The realist may appeal to this idea to answer the perspectivist challenge. While there seems to be something absurd in the notion that the terms and categories of a natural language could be somehow built into the world, the notion that the world instantiates mathematical properties is far less mysterious, at least at first blush. In its contemporary form, mathematization is associated with structuralism. Structuralists hold that inquiry into some domain can reveal only the structural properties of that domain, where structural properties are mathematical properties. Structure is often understood in set-theoretic terms; as Frigg and Votsis (2011) explain,

A structure S consists of (a) a non-empty set U of objects, which form the domain of the structure, and (b) a non-empty indexed set R (i.e. an ordered list) of relations on U , where R can also contain one-place relations (i.e. monadic properties). ... Two structures $S_1 = \langle U_1, R_1 \rangle$ and $S_2 = \langle U_2, R_2 \rangle$ are isomorphic iff there exists a one-to-one (i.e. bijective) mapping $f: U_1 \rightarrow U_2$ such that f preserves the system of relations of the two structures in the following sense: for all relations $r_1 \in R_1$ and $r_2 \in R_2$, the elements a_1, \dots, a_n of U_1 satisfy the relation r_1 iff the corresponding elements $b_1 = f(a_1), \dots, b_n = f(a_n)$

in U_2 satisfy r_2 , where r_1 is the relation in R_1 corresponding to r_2 in R_2 (i.e. have the same index in the indexed sets R_1 and R_2).

Or in slightly simpler terms, for S_1 to have the same structure as S_2 is for S_1 to be isomorphic to S_2 , where S_1 is isomorphic to S_2 iff for every element in U_1 , some element in U_2 corresponds, and when two elements in U_1 bear some relation r_1 , the corresponding elements in U_2 bear corresponding relation r_2 , and vice versa. To illustrate, suppose that S_1 is a set of 10 balls, each of different sizes, and S_2 is a set of 10 sticks, each of different lengths. We can order the balls by the relation “larger than” and the sticks by the relation “longer than”. We can then define a mapping from each ball to each stick, such that for any two balls where one is larger than the other, there are two sticks where one is longer than the other. The set of balls has the same structure as the set of sticks. This abstracts from the nature of the objects and relations: from a structural point of view, it is irrelevant that the objects are balls and sticks, and that the relations are “larger than” and “longer than”.

According to the structuralist, theories present mathematical structures, and these can be mapped to the mathematical structures instantiated in the world. In this sense, there is literally a correspondence between our theories and the world. Scientists speak the language of nature by mathematizing their theories. Nguyen and Frigg (2017) call this the “mapping account”. In the description of structure quoted from Frigg and Votsis above, just think of S_1 as a theoretical structure and S_2 as a structure instantiated by a target system in the world, with f mapping $S_1 \rightarrow S_2$. Note that this account can be weakened, so that instead of taking S_1 to be isomorphic to S_2 , structuralists might favour some other relation such partial isomorphism or homomorphism. The key point is simply that target systems in the world instantiate mathematical structure, and then models and theories can be developed that map to some of this structure. One of the things that’s appealing about structuralism here is that from a structuralist point of view, the nature of the objects that instantiates the structure is irrelevant. A set of balls and a set of sticks can have the same structure. In the same way, then, a system in the world, and a theory or model created by scientists, can have the same structure. This is how we make sense of the idea of the language of nature.

Of course, the mathematization response will be limited to those cases where it is plausible to interpret the products of scientific activity in mathematical terms. I am skeptical that this is as broadly applicable as sometimes seems to be assumed. Many models may be better interpreted as verbal narratives (Dimech 2017) or as collections of propositions (Thomson-Jones 2012). Even if every model can be presented in mathematical language, this often does not make much sense of how some models are used. If we take nonmathematical models seriously, then tension will arise for the structuralist response to failure of fit insofar as we have good reason to take those nonmathematical models to be true or accurate – which in many cases, it seems, we do. This should at least raise the question of whether it is really the mathematical structures that are doing the work of resolving the failure of fit problem, i.e. that have the role of the “language of the world”. However, I will put this concern to one side, since I think there are severe problems for the mathematization response, even ignoring the issue of nonmathematical models.

I do not intend here to criticize structuralist accounts of science. The question is whether appeal to mathematical structures provides a plausible solution to the perspectivist challenge. Can we take mathematics to be the “language of nature”? On this point, there are several serious problems. First, note that the physical world is not in any obvious way a mathematical structure, since mathematical structures are generally taken to be abstract. Abstract objects have no spatiotemporal properties, so the connection between real pendula, which have properties such as length and period, and a model pendulum, a mathematical object that cannot literally have length and period but is used to represent and understand the length and period of real pendula, is somewhat mysterious (cf. Thomson-Jones 2010).

Although in describing structuralism, I spoke of target systems “instantiating” mathematical structure, the relation between physical systems and mathematics is extremely controversial. Part of the problem here is that there is no consensus on the ontology of mathematics in the first place, so we are inquiring about the relation between the physical world and... what exactly? According to fictionalists (Field 1980), eliminative structuralists (Hellman 1989), and Meinongians (Sylvan 2003), mathematical propositions purport to describe abstract objects and structures, but no such entities exist. On these views, mathematics is a useful tool for representing

physical systems, but talk of physical systems “instantiating” mathematical properties is a *façon de parler*. There are no mathematical structures, so there is no relation between them and the world, whether it be identity, or instantiation, or mapping, or whatever else. Now for most purposes in philosophy of science, such problems can be put aside. But for anybody proposing mathematization as a solution to the failure of fit challenge, I don’t think this issue can be avoided, because at issue here is a claim about the world itself, not just a claim about scientific representation (say, that representation involves presenting mathematical structures).

It’s worth noting that even many philosophers of science who defend structuralist approaches are wary of attributing structure to the world directly. The point of structuralism is that a model represents a target system when there is some morphism between the model and the target system. But this is a simplification. We might instead hold that what is required is a morphism between the model and a *data model* of the target system, where the relation between data models and target systems is understood in different terms (e.g. van Fraassen 2008). This move is made precisely because of the problems with taking mathematical structure to be literally instantiated in the world itself.

Moreover, even if we suppose that mathematical structures are instantiated by target systems, using this as a response to the failure of fit argument runs into the problem that we need some prior articulation of the target systems in question. This is quite clear when we consider the definition of structure. As noted, a structure is instantiated by a domain of objects, plus relations on those objects. So, to identify the structure of some system, we need to identify that there are certain objects and relations in the system. This need not involve identifying the nature of those objects and relations – as we saw, it doesn’t matter whether we are dealing with cakes or sticks, or “larger than” or “longer than”. However, the number of objects does make a difference. If U_1 contains ten objects, and U_2 contains eleven objects, then there cannot be a one-to-one mapping from the elements of U_1 to the elements U_2 . Now consider, for example, a biological population. In studying this population, scientists will need to carve it up into individual objects. In principle, the world itself does not compel any particular way of carving up the population: scientists could take individual organisms as the elements of the structure, or pairs of individuals, or the

left sides and right sides of individuals, etc. Putting aside these “artificial” classification systems, it is clear enough that, *in practice*, there are different ways of carving up the population (this point will be elaborated in much more detail in the next chapter). There are different individuality concepts, which will deliver different sets of individuals, and these may be reasonably applicable to the same population in the same theoretical context.

Nguyen and Frigg propose that the application of mathematics to a target system requires the construction of “structure-generating descriptions” of the system (2017: 14). This is a description that specifies the relevant objects, properties, and relations of the target. The scientist must first select those aspects of the target that she is interested in and divide the system up in a particular way. For the biological population, the scientist will specify the members of that population, guided by some individuality concept, and then she will specify certain relations, such as relations of reproduction, parent/offspring, and so on. This is a description of the system in physical terms, not a specification of a mathematical structure. The next step is to schematize the description, which involves abstracting from the specific natures of the objects and relations given in the description. Terms referring to biological organisms are replaced with terms referring to objects. Terms defining relations are replaced with terms that specify only the extensions of those relations. So “parent of”, for instance, will be replaced by a term that tells us which objects the relation holds between, but nothing more. At this point, we have a description of a structure. This structure can then be mapped to a mathematical structure specified by a theory or model.

The application of mathematics involves moving from a physical description D_P to an abstract, structural description D_S . D_S is then mapped to a mathematical structure. Crucially, however, we cannot eliminate D_P . D_P plays an essential role in coordinating the mathematics to the target. Without D_P , we have only abstract structures, which can no longer be taken to apply to anything. So, Nguyen and Frigg say, “a structure applies to a target always only relative to a structure-generating description” (2017: 15) – but structure-generating descriptions are given in natural languages, and so now, we are back to the question of the lexicon-dependence of the world.

It is an obvious point that mathematical formalism in itself says nothing about the world. A person can do mathematics, simply with the goal of exploring a mathematical system. For a mathematical structure to have an empirical application, it must be interpreted. The point here is that this interpretation cannot be given in purely mathematical terms. We must identify aspects of the structure with aspects of the world: we need some sort of mapping from the model to the target. Now, the defender of the mathematization response might be tempted to say that what happens here is that we simply identify the structure of the model with the structure of the world. But what exactly is “the structure of the world”? If this is supposed to be a purely mathematical structure, without physical interpretation, then we are no longer talking about the world. This would just be a matter of presenting two mathematical structures. There is a difference between presenting a structure for its own sake and presenting a structure as the structure of some worldly system. Mathematics alone is not sufficient for the latter.

This point arises in the debates around structural realism. Structural realists need a criterion that distinguishes structure and content, in a way that picks out the features of successful past theories that were retained, and crucially, that can be taken to describe *unobservable* structure. As Stanford (2006: 181) notes, Worrall, in his original presentation of structural realism, (Worrall 1989) sometimes talks as if structural content is simply a matter of the mathematical formalism, so that we are justified in believing nothing more than the equations and other abstract structures of our theories. The problem with this is that it's simply not enough to preserve belief in the mathematics, because the mathematics can be satisfied by empirical, observable structures alone. Take, for instance, Newton's laws. *Of course* Newton's laws can be (approximately) recovered in later science: these laws approximately accommodate observable regularities. If a theory has successfully captured observable regularities, and a later theory also captures those observable regularities, then there must be some mathematical structure retained in the transition from the earlier to the later theory, no matter how significant the underlying theoretical changes are. The moral is, as in the previous discussion: the mathematics needs to be interpreted via structure-generating descriptions, and any given mathematical formalism can be true of many things. Without a structure-generating description, it is not true of *anything*.

I have framed these arguments in terms of the structuralist approach. Could the defender of the mathematization response think of mathematization in different terms, so as to avoid the problem? This is doubtful. For one thing, this problem would, I suspect, arise for any plausible conception of mathematical properties, because what is really driving this problem is that mathematical objects and relations are abstract. The difficulty is how we coordinate abstracts with physical systems. Second, the structuralist approach has the significant virtue of providing a clear, well-developed account of what exactly “mathematical language” is, which is applicable to mathematical models across the board. Many different mathematical representations are available for what the structuralist would treat as the same structure. As a result, if we abandon the structuralist approach, we only increase the abundance of mathematical representations of any given target system. Consider Galileo’s mathematization of free fall. Today, relations that are represented algebraically were originally presented by Galileo using geometrical representations (Forinash et al 2000). Galileo understood free fall in geometrical terms, of lines, points, and proportions. Indeed, the full quote which I noted as inspiration for the mathematization response is that the book of nature is “written in mathematical language, and its characters are triangles, circles and other geometric figures...” The algebraic representation and geometrical representational are equally accurate, though today we take algebraic representation to have greater pragmatic virtues. But which is the language of nature? The structuralist avoids this problem, by taking both representations as describing the same structure.

There is one final difficulty for the mathematization response. This is that, as argued by Cartwright (1983) and Giere (1988), among others, the mathematical structures used in science often do not even *prima facie* map to real-world systems, but only to idealized models of those systems. I noted above the concern that if the ideal pendulum is a mathematical object, then it cannot have spatiotemporal properties. But even if we suppose that the ideal pendulum does somehow have spatiotemporal properties, it clearly does not have the same spatiotemporal properties as any real pendulum. And the differential equation that describes the motion of an ideal simple pendulum is not true for the motion of any real pendulum. Now, some of the idealizations in the model pendulum can be removed, and this will increase the

complexity of the equations describing its motion. It is far beyond our computational capacities to remove all idealizations, however.

In practice then, mathematical structure is not taken to map to the world itself. The mathematization of science instead proceeds via the introduction of idealizing assumptions. There is much to be said about why exactly idealizations are useful, and how it is that we can use falsehoods to learn truths about the world. The point in this context is simply this: the fact that mathematical structures are useful devices for scientific representation does not show that mathematical structures objectively correspond to the world, because there are a host of cases – just pick any idealized mathematical model – in which we know that these structures are useful but do not correspond. In this respect, mathematical description seems to exhibit much the same problem as natural language description. It is precisely this point, as we saw, that drove Giere to adopt a similarity account of representation. Of course, it may be argued that a real pendulum could *in principle* be described in perfect detail mathematically, that there is some equation which truly describes its behaviour, even if we could never discover this. But similarly, there is just as much reason to suppose that there is a natural language description of the pendulum that is completely detailed and accurate.

None of this is to say that structuralism in general, or even structural realism in particular, is mistaken. The question is whether mathematization and structuralism can support the notion of a “language of nature”, a mind-independent correspondence between our theories and the world. There are various reasons for skepticism on this point.

6.7. *Natural kinds*

The most popular response to the failure of fit argument is no doubt the appeal to natural kinds. According to standard realism, we have good reason to believe that scientific classification schemes match the natural kinds. That is, the world is divided into kinds – the lexicon of the world is a lexicon of kinds – and this division is mirrored by the classifications generated by our best theories. Furthermore, scientific practice is in this respect just a more sophisticated development of basic capacities

found across the animal kingdom. There is in general good reason to believe that our representations “fit” the kinds of the world. In the next chapter, I will argue that attention to how scientific classification schemes are developed and applied undermines the picture of scientific classifications as mirroring the divisions of nature. I devote a separate chapter to this discussion partly because it requires careful development, and partly because my goal is not simply to attack realist approaches to kinds, but to present a perspectivist alternative. After all, even if traditional realist conceptions of a language of nature fail, this would not in itself constitute a convincing case for perspectivism. The failure of the realist defences explored in this chapter can only be one part of the perspectivist argument. For all that has been said so far, we might simply remain agnostic between traditional realism and perspectival realism. My discussion of kinds in the next section will flesh out the failure of fit argument, in particular, the importance of simplifying in response to complexity.

Chapter 7

Natural Kinds

7.1 Introduction

In the previous chapter, I examined several initial responses to the failure of fit argument and found that none of them were successful. There remain, however, however responses to consider. The first is the appeal to natural kinds. Most realists endorse, in one way or another, the claim that the world contains natural kinds and that our best theories succeed in tracking those kinds. Indeed, this is a popular enough position, and has been developed with enough sophistication, that the failure of fit argument might seem hopelessly implausible. Recall Kuhn's statement of the central idea of the argument:

Evaluation of a statement's truth values is...an activity that can be conducted only with a lexicon already in place, and its outcome depends on that lexicon. If, as standard forms of realism suppose, a statement's being true or false depends simply on whether or not it corresponds to the real world – independent of time, language, and culture – then the world itself must be somehow lexicon-dependent.

The realist can respond that the "lexicon" of the world is a lexicon of kinds. There are objective divisions in nature. Our representations mirror those divisions. This is the sense in which there is an appropriate correspondence between our theories and the world. So one task of the next two chapters is to undermine realist approaches to kinds. The second task is to develop in some detail a perspectivist alternative. This is important because the failure of fit argument, as stated so far, has the form of a challenge to realism. For all that has been said, we might favour outright antirealism, rather than perspectivism, in response to failure of fit. Or it might be that *nobody* has a convincing solution to failure of fit, in which case, the realist's failure to deal with the challenge would not even count against her position in comparison with other available views. The development and defence of a perspectivist account of kinds will complete the case for perspectivism over traditional realism. Additionally, we will see that the perspectivist account of kinds provides a novel approach to the old topic

of classification and kinds, and helps to resolve some of the central problems in that debate.

7.2. Traditional approaches to classification

What is the importance of classification? On the one hand, classification can seem a mere matter of semantics. What does it matter what we call things? Yet on the other, all scientific theories rely on drawing distinctions between different kinds of things. Theorizing cannot even get started without some initial classification scheme for the phenomena of interest, and as theories develop, classification schemes change to better express theoretical results. A general assumption is that it is the aim of science to produce “natural” classification schemes – that is, schemes that “carve nature at its joints”. It seems intuitively obvious that some classification schemes are more natural than others: some reflect the genuine order and structure of the world, while others are artificial, or arbitrary, or dependent on our interests, such that they tell us more about us than about the mind-independent world. Much of the debate about classification is conducted in terms of “natural kinds”, where natural kinds are the kinds in the world that correspond to the natural classification schemes. Do natural kinds exist, and if so, what is their nature?

Dupre (1993: 17-18) suggests three desiderata for a natural classification scheme. First, the classification should exhibit categorical distinctness, that is, it should have sharp distinctions. There should be a determinate fact into which class a particular item falls. If there is instead only continuous variation, then it is up to us where the draw the lines. As Ellis (2001: 20) puts it, “if we have to draw a line anywhere, then it becomes *our* distinction, not nature’s.” We cannot be carving nature at its joints if there are no joints at which to carve. Dupre’s second criterion is that natural classification must – in some sense – be discovered, not merely invented. Of course, we have to decide to construct a classification scheme in the first place. But the distinctions in the scheme must correspond – again, in some sense – to distinctions in the world. The final desideratum is that assigning an object to a particular class must allow us to reliably infer as many other properties of the object as possible. The ideal would be that the classification tells us everything about the object. Dupre’s desiderata usefully summarizes the traditional realist view of natural kinds. It is a

view with a long history; see for instance Whewell (1860), who suggests that scientific investigation reveals “classes in which we have not a finite but an *inexhaustible* body of resemblances amongst individuals, and groups made by nature, not by mere definition.”

Belief in natural kinds is common across the many varieties of realism. Votsis (2012: 86-87) suggests that entity realists, structural realists, semirealists, and more, are committed to the success-to-truth inference, and to the view that successful theories are true partly in virtue of tracking the objective natural-kind structure of the world. It is easy to see the appeal of this traditional realism. It can seem just obvious that there are real, substantive divisions in nature, even before we consider sophisticated scientific theorizing. Look at the night sky: black with dots of white; surely this is a clear division. With more theoretical understanding, the division becomes even more obvious: large areas of mostly empty space, with relatively tiny points of large mass, shining brightly. This is a real division in the world, a difference, not something imposed by us. But there are two points worth bearing in mind here, which help to clarify the alternative, conventionalist position.

First, in this case, it seems that the *obvious* division is in part imposed by our unique biological perspective. We look at the night sky and distinguish stars from empty space. There are conceivable perspectives from which the kind *star* would be rather less obvious. From an intergalactic point of view, if we are considering the development of whole galaxies over time, perhaps there is no need to distinguish individual stars, or even distinguish stars in general from clouds of gas and dust, and so on. Second, the *obvious* divisions are often not particularly useful in scientific theorizing. Hence “star” in the colloquial sense arguably has a different meaning from “star” in the technical sense. The technical use demarcates stars in terms of the process of thermonuclear fusion occurring in the core, distinguishes stars from brown dwarfs, etc.; people talked about stars long before they had any understanding of such things. Similarly, we can divide the universe into terrestrial and extraterrestrial phenomena. This is a real distinction, as real as the distinction between stars and planets; and what’s more, it’s a distinction of obvious importance to us in our everyday lives. But sorting entities in terms of whether they are terrestrial

or extraterrestrial does not reflect natural kinds. Earthly objects are privileged to us because we live on Earth.

Why should we believe in natural kinds? How do we go about constructing natural classification schemes, and how do we justify the claim that a given classification scheme is natural? Within philosophy of science, there is a consensus in favour of a naturalistic approach. The assumption is that if anything provides access to the natural kinds, it is science; if natural classification schemes are to be found anywhere, it is within our best scientific theories. Our philosophical account of natural kinds must accommodate the results of the sciences. This is supported by the standard realist epistemology. As noted in the introduction, contemporary realists appeal to inference to the best explanation to motivate ontological commitment. We are justified in taking our best theories to provide true descriptions of the world – or more precisely, in this case, we are justified in taking them to provide classification schemes that carve nature at its joints – because that is the best explanation for the predictive, explanatory, and manipulative successes of those theories. Since this point will be important to my overall argument, it is worth noting some representative examples of this approach.

Boyd (1999: 146):

It is a truism that the philosophical theory of *natural* kinds is about how classificatory schemes come to contribute to the epistemic reliability of inductive and explanatory practices ... [T]he theory of natural kinds is about how schemes of classification contribute to the formulation and identification of projectible hypotheses.

Chakravartty (2007: 152):

The primary motivation for thinking that there are such things as natural kinds is the idea that carving nature according to its own divisions yields groups of objects that are capable of supporting successful inductive generalizations and predictions. So the story goes, one's recognition of

natural categories facilitates these practices, and thus furnishes an excellent explanation for their success.

Koslicki (2008: 790) suggests that those who are realist about natural kinds:

are typically motivated in their belief in the existence of natural kinds by the role these kinds play in (i) induction and prediction; (ii) the laws of nature; and (iii) causal explanation.

Taylor (2020: 2) says that in investigating natural kinds:

we aim to discover which categories pick out genuine scientifically interesting divisions in nature, worthy of investigation. This is opposed to the categories that pick out groups of entities that have no scientific interest or import. Identifying natural kinds and disregarding categories that are not natural kinds is a key source of progress in science as a whole.

The general assumption underlying realistic, naturalistic approaches is that natural kinds relate to the inductive and explanatory success of science. Our best scientific theories track the natural kinds, and this explains their successes. So, our account of natural kinds should be based not on *a priori* grounds such as thought experiment, conceptual analysis, stipulative definition, etc., but rather on the results of science. Of course, this need not involve the assumption that *all* the kinds found in scientific theories correspond to natural kinds. It is a familiar point that many theories have been discarded, and with their disappearance, many theoretical terms that had been taken to refer slipped out of use. Nobody supposes that “caloric” or “phlogiston” name natural kinds. We also know that current theories are incomplete, and likely to change at least somewhat in the future. However, I take it that this point can be dealt with by the usual appeal to “approximate truth”. The realist will say, more precisely, our theories are not true but only approximately true; similarly, our theories do not exactly track the natural kinds but only approximately track the natural kinds. We are mistaken about some kinds, and we are ignorant of others; but by and large, our best theories get it right. It is the task of philosophy to provide a theory of natural kinds that makes sense of this and explains how this can be the case.

Following Franklin-Hall (2015), I will call this general assumption *Coordination*. Scientific classification schemes and natural kinds must in some sense “match” one another; they must be coordinated with one another. This forms the response to the failure of fit challenge. But this also gives rise to the *Coordination Problem*: the task of showing how scientific kinds could be coordinated with natural kinds. Now obviously, there is no guarantee that we will ever in fact achieve Coordination; we cannot conclusively rule out that scientific classification schemes simply fail to latch onto the joints of nature. We might, however, provide a philosophical theory that accommodates Coordination and explains how Coordination could be achieved, given the assumption that our best theories are broadly accurate. The Coordination Problem is the problem of showing that a realist interpretation of natural kinds is coherent and compatible with the practice of science. We can compare the goal here to that of externalist responses to skepticism about the external world. The externalist cannot *guarantee* that there really is an external world; even under some externalist account of knowledge such as reliabilism, it is logically possible that I am a brain in a vat. But she can show how, given that there is an external world, I could come to have knowledge of it. This is the best we can expect in response to the skeptical challenge. Along the same lines, we cannot demand that the realist guarantee that Coordination does occur. But her position should be coherent by its own lights. So, we should be able to see how the truth of our best theories would ground Coordination.

With that said, let's turn to some of the realist approaches to natural kinds in the philosophical literature. This must necessarily be brief; I intend simply to highlight some general tendencies in the debate. First then, one of the oldest realist positions is essentialism. According to essentialism, kinds are distinguished by essences, i.e. some intrinsic property or set of intrinsic properties whose possession is necessary and sufficient for membership of the kind. The essence of a kind governs how the kind interacts with other things. For example, the essence of gold is 79 protons; all and only atoms of gold have 79 protons, and the fact that these atoms have 79 protons explain the behaviour of gold in various circumstances. Sober (1980) treats essentialism as committed to three claims: (1) Certain properties are necessary and sufficient for kind membership, so are shared by all and only the members of the

kind. This set of properties is the essence of the kind. For instance: all and only helium atoms have two protons. (2) The essence must be explanatory; it must explain other properties of the kind, and how the members of the kind interact with members of other kinds. Thus, the fact that helium atoms have two protons explains the other properties and interactions of helium atoms. The diversity of all the different forms of helium is explained by how this *shared* underlying structure manifests different properties in different contexts (in different pressures, different temperatures, etc.). (3) Kinds are defined in terms of their members – Sober calls this a “constituent definition”. Once we understand the properties of the kind-members, we understand the properties of the kind itself.

There are, of course, different forms of essentialism. Sober’s criteria provide a representative example. At the very least, Sober’s first criterion – that kinds are distinguished by necessary and sufficient intrinsic properties – is a basic requirement of an essentialist position. It is also the source of the most trouble for essentialists. How does essentialism fare with respect to the Coordination Problem? Well, if our best theories track essence kinds, this goes some way to explaining the predictive and explanatory successes of those theories. It’s easy to see how knowledge of essences would support us in formulating projectible hypotheses, laws of nature, causal explanations, and so on. Unfortunately, the simple fact is that outside of physics and chemistry, very few scientific theories postulate essence kinds. Species are paradigmatic examples of natural kinds, but species do not have essences. The organisms of a given species will tend to share various properties across a wide range of environments, due to common causes (common ancestry, reproductive isolation, developmental constraints, occupation of a similar niche), but in biology, variation is the norm, and no property or set of properties will be necessary or sufficient for species membership. Such examples leaves the natural kind realist with a dilemma (cf. Chakravartty 2007: 158). Either natural kinds are of very little relevance to scientific practice in many cases – in which case, we would have to give up on Coordination – or a more liberal conception of natural kindness is required. In many domains, scientists just don’t work with essence kinds, so if we want our account of natural kinds to be naturalistic, we need to weaken the conditions for natural kinds.

This has led to the development of *cluster kinds*, which is the consensus account of kinds in the philosophy of the special sciences. The most prominent cluster kind account is Boyd's homeostatic property cluster (HPC) account (see e.g. Boyd 1999; 2010; 2019). According to the HPC view, certain sets of properties tend to cluster in nature, in that these properties reliably co-occur across a wide range of conditions, where the co-occurrence of these properties is due to what Boyd describes as homeostasis. This involves either (a) some of the properties in the cluster tend to favour the occurrence of other properties in the cluster or (b) the properties in the cluster arise due to some shared underlying causal mechanism, or both. As a result, deviations from the cluster are relatively unlikely to occur, at least in particular specified conditions: if some of the cluster properties are occurring together, you are likely to find many of the others as well. HPC was explicitly designed to capture cases for which essentialism seems implausible, in particular, biological species (Boyd 1999). Natural kinds are stable clusters of properties. Some property clusters may be very tightly connected, others less so.

How does this play out with respect to Coordination? As the quote above demonstrates, Boyd is explicit that our commitment to HPC kinds is secured by their use in successful practices of induction and explanation. More precisely, Boyd argues that classification schemes pick out natural kinds when there is *accommodation* between our representational practices and the inductively or explanatorily relevant causal structures of the world (Boyd 2019: 6). Accommodation in the sciences is simply a much more sophisticated form of an epistemic achievement can be found throughout the animal kingdom. Boyd illustrates this with an example of ground squirrels. Ground squirrels have two types of alarm calls, one for aerial predators and one for terrestrial predators. The aerial calls tend to be produced in response to aerial predators, rather than in response to other features of the environment relevant to survival, and hearing the call will initiate a particular pattern of behaviour in the squirrel, appropriate to avoiding aerial predators. Similarly for the terrestrial calls. So the signalling system and social structures of the squirrels "are *accommodated* to relevant causal features of their environment so as to facilitate predator avoidance" (Boyd 2019: 7). In the human case, certain classification schemes help to achieve accommodation of human inferential practices to causal structures in the world, so as to facilitate predictive and

explanatory success. For squirrels to achieve predator avoidance, they need a signalling system. So it is with scientists: for their inductive and explanatory generalizations to be accommodated to the world's causal structures, they need a vocabulary with which to express such generalizations; this vocabulary must itself be accommodated to the causal structures. This is just to say: the vocabulary must pick out the natural kinds.

Boyd's current definition of HPC kinds is as follows (Boyd 2019: 30):

K is an HPC kind in a disciplinary matrix, D, just in case

(1) K is a natural kind in D in the accommodationist sense, and

(2) There is a (typically imperfectly) homeostatically buffered temporally and spatially extended causal process (a clustering), C (type or token, depending on the kind), such that participation in C defines membership in K. Reference to K contributes to accommodation because of the inductively/explanatorily relevant causal profiles of Ks are those that are (typically imperfectly) homeostatically stabilized by their participation in C.

Notice that kinds are defined relative to disciplinary matrices. This creates a role for human interests in fixing what kinds there are. Each of our practices exhibits certain achievements. This can only occur when that practice is accommodated to causal structures. So what counts as a HPC kind will depend on what exactly our goals are in conducting a particular inquiry, and how exactly we go about achieving them.

Boyd provides a sophisticated account of Coordination, in the context of a more liberal approach to kinds than traditional essentialism. But HPC may not be liberal enough. According to promiscuous realism (Dupre 1993), there are many equally legitimate ways of dividing the world into kinds, and these different classifications schemes may be cross-cutting. Within biology, we find various criteria for species membership. For example, the biological species concept defines a species as any group of organisms capable of interbreeding, that are reproductively isolated from other groups. The ecological species concept defines a species as any group of organisms that occupy the same ecological niche. Phylogenetic species concepts define species in such a way as to rule out polyphyletic groups, that is, groups of

organisms that do not share an immediate common ancestor. These species concepts carve up the world in different ways – applying the biological species concept will produce a different taxonomy than applying a phylogenetic species concept. The promiscuous realist embraces pluralism here, and holds that both of these species concepts, and others, accurately capture some of the mechanisms that produce biological diversity. Which classification scheme we should use is dependent on our interests. There are many contexts where we need to classify organisms, but where the biological species concept is not operational since it would be difficult or even impossible to determine relations of interbreeding, such as a study of a speciation event that took place in the past. Promiscuous realism is therefore committed to pluralism, and to the view that even though scientific representations map onto mind-independent features of the world, they are also moulded by conventions, historical contingencies, and the interests of particular research groups. “[W]hat should be the grounds for accepting a taxonomic scheme ... [are] not that it is the right one, since there is none such; but that it serves some significant purpose better than the available alternatives” (Dupre 1993: 51-52).

Furthermore, for promiscuous realists, “significant purpose” does not necessarily mean “scientific purpose”. Promiscuous realism allows that not just scientific classification schemes, but the classifications produced by a wide variety of practices and inquiries may track natural kinds. The commonsense kind “lily” refers to a type of flower. The biological family Liliaceae includes various species that are not traditionally classed as lilies, including onions and garlics, while excluding various species that are traditionally classed as lilies, such as canna lilies and arum lilies. There is a commonsense kind *lily* and a biological kind *Liliaceae* which provide cross-cutting classifications. But in both cases, the kind is tracking certain properties in the world. The question is simply which classification is appropriate given our interests. It is worth noting that more current versions of the HPC view have been liberalized in a similar fashion. Boyd’s current position is that different practices may accommodate to different causal structures, producing cross-cutting classifications; he cites “lily” as a natural kind that is relevant to the practices of gardening and landscaping (Boyd 2019: 35).

A similar approach is Chakravartty's account of kinds as "sociable properties" (Chakravartty 2007: 168-180). Properties are not randomly distributed. Rather, they tend to group together. The strongest degree of association is found in essence kinds, where sets of properties always co-occur. Then there are the looser associations of properties in the cluster kinds that motivate the HPC account and promiscuous realism. For Chakravartty, however, literally any shared property, any resemblance, even if it is only a loose association, counts as natural kind. We exploit only a tiny fraction of the innumerable natural kinds in our everyday and scientific classification schemes. "[I]t is the fact that members of kinds share properties, to whatever degree, that underwrites the inductive generalizations and predictions to which the categories lend themselves" (Chakravartty 2007: 170).

If natural classification schemes need not be scientific, does this involve a repudiation of Coordination? Not really. It is just that what counts as a "successful practice" is expanded to contexts beyond science. It's quite natural to extend our explanation for the success of science to successes in other fields, especially if, like Dupre (1993), we reject the unity of science thesis that science operates with a single method. In any case, even on Chakravartty's view, it remains the case that our best scientific theories coordinate with the world. It's just that various other, non-scientific practices also achieve Coordination. What matters for Coordination is simply that our best scientific theories track the natural kinds. If other practices are successful in such a way as to track natural kinds, so much the better.

7.3. Perspectivist accounts of kinds

What role is left for perspectivism? Perspectivism about kinds is discussed at length by Cretu (2020). She begins with what she calls the "commitment problem" for perspectivism. Perspectivism involves, Cretu argues, a realist commitment and a perspectival commitment. The realist commitment is that science delivers genuine knowledge of the world; the perspectival commitment is that this knowledge is always situated within a particular context: it is perspectival knowledge. The commitment problem is the question of how to integrate these two commitments coherently. Cretu uses this problem as a springboard for her own perspectivist account of kinds, which will be discussed in detail shortly. However, before

examining Cretu's own position, it's worth noting a problem with her line of argument here. It seems to me that the commitment problem has already been solved. We appear to have a number of options for integrating what Cretu calls the realist commitment and the perspectival commitment.

Cretu describes the HPC account as a "strong realist" view which is not sufficiently perspectival. Why not? For one thing, she argues, because it is not sufficiently practice-relative. Now, perhaps this is correct for early versions for the HPC account. But as we have seen, Boyd explicitly builds practice-relativity into his current theory. Indeed, Boyd even states that "definitions of particular natural kinds are relative to human projects and disciplines ... on the accommodationist conception, natural kinds are mind-dependent social constructions" (2019: 9). It is still a form of realism because these social constructions are accommodated to causal structures. But natural kinds arise from the interaction between our epistemic practices and the causal structures in the world.

The more liberal accounts of kinds, such as promiscuous realism, will make even more room for practice-relativity; many of them are explicitly designed to capture the ways in which our classifications are dependent on our interests. Indeed, Chakravartty explicitly rejects the standard objective/subjective distinction as useful with respect to natural kinds, saying that while classification schemes track objective similarity relations, "this offers no prescription ... against utility, convenience, and interests playing a role in the demarcation of kinds" (2007: 175). Thus, the views on the table already address Cretu's commitment problem. What's more, the reasons why these views make room for the influence of human interests and conventions on our classification scheme is broadly the same as the reason why the perspectivist appeals to these factors: the enormous complexity of the world and our need to engage in abstraction and simplification; our inability to completely step outside our background goals and assumptions and achieve a "view from nowhere", etc.

Indeed, if anything, perspectivist proposals seem to be committed to a stronger form of realism than the more liberal realist views. A good example of this can be found in Massimi's (2017) outline of a perspectivist view of kinds. She contrasts perspectivism to other forms of pluralism, in particular promiscuous realism. On

Massimi's view, the perspectivist is not committed to pluralism with respect to taxonomic classifications, but rather a kind of epistemic pluralism. This is understandable given, as we have previously seen, Massimi's treatment of "perspectivism" as a thesis not about what we have scientific knowledge of, but rather about *how* we form scientific knowledge. The perspectivist, Massimi says, should hold that there are mind-independent natural kinds, but that our ability to know what the natural kinds are is dependent on our perspectives. She contrasts this with Dupre's view, which she takes as being non-committal with respect to the mind-independence of kinds. As Massimi puts it: "Our multiple *ways of knowing* do not track multiple mind-independent kinds in nature" – so perspectivists should resist ontological pluralism.¹

So, on the one hand, perspectivism seems very much at home with many of the contemporary approaches to natural kinds. The good news is that this provides more evidence that perspectivism is perfectly coherent, at least in its basic commitments. The bad news is that this highlights a serious problem, which I will call the *Redundancy Problem*. It just isn't clear that perspectivism has anything to bring to the debate about natural kinds that hasn't been said before. What is often treated as the unique contribution of perspectivism – its emphasis on the way that knowledge involves an interaction between scientists and the world, its commitment to both realism on the one hand, and to the cultural and historical situatedness of theories on the others – appears to have been developed in some detail already by realist positions in the natural kinds debate. It seems as if perspectivism is already a feature of the standard positions!

¹ I am not convinced that promiscuous realism is non-committal about mind-independence in the way Massimi suggests. I think what motivates the promiscuous realist is rather the concern that there are, in a sense, too many real kinds. Whether or not similarity relations hold between sets of objects is a mind-independent fact. Any similarity relations might be the basis for a classification scheme. But it would be unwieldy to even attempt to track them all; we must find some way of bringing order to the chaos and *the world itself* cannot tell us how to do this, so we use our interests and conventions to guide us. So it is not clear that Massimi presents an alternative to a promiscuous realist account of kinds.

In principle, the perspectivist might simply accept this, and argue that the distinctive contribution of perspectivism lies elsewhere. But it's worth noting that without a solution to the Redundancy Problem, the perspectivist faces a dilemma. On the one hand, if realists can appeal to natural kinds in a way that is compatible with perspectivism, then it seems that realists have a straightforward solution to the failure of fit challenge. The lexicon of the world is a lexicon of kinds, and our best theories provide classification schemes that match those kinds. This is where there is a "fit" between our theories and the world. So now the central argument for perspectivism is undermined. Not only does perspectivism fail to make a distinctive contribution to the natural kinds debate, but there is no longer any difficulty for realism, and so no longer any reason to take perspectivism seriously in the first place. Alternatively, we might argue that the appeal to natural kinds does not resolve the failure of fit challenge. But now the worry is that, if all realist positions are challenged by failure of fit, and if perspectivism has the same account of natural kinds as many of these realist positions, then failure of fit generalizes to perspectivism as well. The perspectivist is claiming, like the realist, a fit correspondence between our classification schemes and the world itself. In this respect, the perspectivist faces the same challenge as the realist in treating the world as somehow lexicon-dependent.

Let's recap the general development of the debate. There is a consensus in favour of Coordination: the classification schemes of our inductively and explanatorily successful practices map onto the joints of nature. We take it that our best scientific theories track the natural kinds (and perhaps the theories and practices of other domains also track natural kinds, but certainly our best scientific theories do). Essentialism is far too strict, since it appears to be applicable only to some of the kinds distinguished in the hard sciences. So the concept of kind has been progressively weakened, to the point that on some concepts of kind, just any shared set of properties is sufficient to count as a natural kind. We also see an increasing acknowledgement of the human contribution to classification. There are two problems that will be in the background for the remainder of this chapter and the next. Concerning the natural kinds debate, we have:

The Coordination Problem: How is it possible for scientific kinds to be coordinated with natural kinds?

For a perspectivist account of kinds, we have:

The Redundancy Problem: What new contribution does perspectivism make to the natural kinds debate?

I will now turn to perspectivist approaches to kinds, focusing on the proposal given by Cretu (2020). I will argue that Cretu's account fails, and that it falls to the same kind of problems that afflict earlier theories of kinds. However, it also contains important insights that point the way to a more promising perspectivist approach to scientific classification. That alternative approach will be developed in Part IV.

7.4. Cretu's perspectivist account

Cretu provides a more detailed perspectivist approach to kinds. The key move in her account is to distinguish mind-independence from perspective-independence. Complete mind-independence for our classification schemes is unobtainable, since all inquiry requires us to make various theoretical and methodological assumptions. But we can achieve perspective-independence. Cretu here draws a distinction between research traditions and perspectives. She follows Laudan (1977: 97) in defining a research tradition as "a set of assumptions: assumptions about the basic kinds of entities in the world, assumptions about how those entities interact, assumptions about the proper methods to use for constructing and testing theories about those entities." A perspective, by contrast, is a specific theoretical framework, involving the background theoretical knowledge and set of theoretical goals that a particular scientist or research group may hold. It is "an evolving theoretical outlook that offers a particularised ontology" (Cretu 2020: 7), and which is applied to interpret the results of the research tradition. Cretu's idea is that natural kinds can be "authenticated" within a research tradition, then various perspectives may be developed to study them; if any given perspective should fail, as they often do, then the kinds remain authenticated on the basis of the broader research tradition. Perspectives offer interpretations of natural kinds that are authenticated

independently, and so the kinds are retained throughout changes in perspective. Thus, “the ontological commitments warranted by the research tradition can be retained even when perspectives are abandoned” (2020: 4). Then a classification scheme can be said to be natural provided it tracks these perspective-independent entities.

How exactly does authentication take place? According to Cretu, what is revealed in research traditions are “real patterns”, in something like the sense of Dennett (1991) and Ladyman and Ross (2007). Real patterns are stable empirical regularities, and these can be tracked by successful theories. Cretu is somewhat unclear about what precisely is involved in authenticating a real pattern, but she does note that the point of authentication is the elimination of experimental error and inferences based on freak results (Cretu 2020: 13). I take it that authentication involves deploying procedures for calibrating instruments, control of confounding variables, various methods of statistical analysis, etc., that allows us to be confident that we have detected an empirical regularity, independently of any claims about the nature of that regularity.

Cretu gives several illustrations of this picture. First, the discovery of positrons. Scientists identified positrons through their observations of things such as cloud chamber photographs, given their general assumptions about particles and about the operation of cloud chambers. Then they developed hole-theoretic and field-theoretic perspectives on positrons. After the hole-theoretic perspective was displaced, positrons remained in our ontology. Similarly, Cretu cites the development of stellar classification: astronomers began by measuring and systematizing the spectra of stars. Stars were classified on the basis of the colour of the starlight and the absorption lines in the spectra. The early period of stellar classification culminated in 1901 with the Harvard classification system with seven spectral classes, from blue to red: O, B, A, F, G, K, M. This classification was authenticated independently of any assumptions about the nature of the stars; for example, it was not known at the time that this was a temperature sequence from hot to cold. We have since developed various perspectives on the origin, evolution, and constitution of stars, but the basic Harvard sequence is, with various additions, still in use today.

So Cretu proposes a two-stage picture of scientific theory development. First, scientists must establish that a real pattern is genuine, and this justifies us in taking a particular classification scheme to correspond to reality. The authentication of a pattern takes place within a broader research tradition, and we are required to designate particular entities and judge the relations of likeness and difference between them, so this is not fully mind-independent. Second, the pattern is then studied further by perspectives developed within the research tradition. If the pattern is properly authenticated, it will be retained through changes in perspective. Natural kinds are simply real patterns.

As Cretu recognizes, this threatens to lead to a huge abundance of natural kinds. After all, there is an infinitude of patterns that can be discerned in the world. If a natural kind is just any real pattern, then tables and chairs turn out to be natural kinds also. Indeed, any set of things that share any properties will be natural kinds, since insofar as they share a property, they can be treated as exhibiting a pattern. Cretu solves this problem by restricting natural kinds to those real patterns that have high indexical redundancy, following Ladyman and Ross (2007: 297). Indexical redundancy is a matter of *measurability*: a pattern has high indexical redundancy when it can be measured in multiple ways, from multiple places, and using multiple instruments. Electrons, for instance, are measurable throughout the universe, and within many frames of reference. Biological patterns have lower indexical redundancy, since they are measurable only in very specific regions. Tables and chairs have very low indexical redundancy, or so Cretu says. The natural kinds are those real patterns with high indexical redundancy.

I will now outline two problems for Cretu's view. First, there is a problem of indeterminacy. Notice that attempting to apply Cretu's position to specific cases of scientific classification is likely to recapitulate the debate that we have already seen in earlier discussions of classification. We want our theory of natural kinds to track the scientific kinds; whatever natural kinds are, it must turn out that our best theories capture the natural kinds. Cretu clearly shares this assumption, insofar as the whole point of introducing indexical redundancy is to exclude various kinds distinguished in non-scientific contexts. But the concern is that if the standards of indexical redundancy are high enough to exclude "real patterns" such as tables and chairs

from being natural kinds, we are also likely to lose the kinds distinguished in many sciences. We must draw the line somewhere on exactly how high a degree of indexical redundancy is required; there is no non-arbitrary way to do this, and there is no reason to expect any particular line to capture the right cases.

This same problem of indeterminacy arises in a different way. I draw here on McAllister (1997). Any given data set is compatible with an infinite number of patterns, because it can be interpreted as exhibiting a pattern plus a particular noise level. Before we can say what the pattern is exhibited in a data set, we must specify the noise level, and this is a matter of stipulation, not something determined by nature. McAllister gives the example of planetary orbits (1997: 226). For Galileo, planetary orbits are circles; for Kepler, they are ellipses; for Newton, they are curves that deviate slightly from ellipses, due to gravitational interaction with other bodies in the solar system. Both patterns may be discerned in the same data set; it's just that with Galileo's pattern, we stipulate a higher noise level. Now, McAllister is clearly defining "pattern" in an extremely fine-grained way. Presumably, Cretu would want to say that Galileo, Kepler, and Newton were all at least approximately correct about the planetary orbits; and so in this sense, they are all investigating the same pattern. In general, then, what counts as the same pattern? At the very least, we must fix a level of grain. If we demarcate patterns in a way that is too fine-grained, it will be impossible to maintain that patterns are retained over research traditions.

The second problem for Cretu is that due to the perspectivist aspect of her view, it seems that she is unable to secure commitment to any natural kinds whatsoever. On the one hand, Cretu wants the results of the research tradition to justify ontological commitment to natural kinds. As noted, Cretu is somewhat ambiguous about the precise distinction between research traditions and perspectives. Still, this distinction is similar to others in the literature. Given that research traditions involve ontological commitments, we might take them as including something like Hacking's (1983: 365) "home truths" about unobservable entities. We are justified in believing in electrons, because there are a variety of well-understood causal properties of electrons that we can deploy in building instruments and investigating nature, and these "home truths" survive changes in the higher-level theoretical understanding of electrons. This is also similar to the view of scientific practice in Bogen and Woodward's (1988)

data/phenomena distinction. Scientists infer the existence of phenomena from data, on the basis of a basic understanding of instruments plus techniques for data analysis, largely independently of higher-level theory. Phenomena arise from a small number of causal factors and occur in a wide variety of different contexts, so are measurable in many ways. “Real patterns” are similar to phenomena in Bogen and Woodward’s sense. Indeed, I am not sure what the distinction would be between saying e.g. that the melting point of lead is a phenomenon vs. the melting point of lead is a real pattern. These analogies bode well for the realist element of Cretu’s perspectivism: on this kind of view, real patterns are mind-independent things in the world, and the development of research traditions justifies belief in particular real patterns. So at first sight, we seem to secure the realist commitment.

The problem is that whether this is a plausible view of *natural kinds* is dependent on exactly how much is contributed by the perspective to our understanding of a real pattern. Assuming we accept the distinction between research traditions and perspectives, are we justified in treating real patterns as analogous to Hacking’s unobservable entities or Bogen and Woodward’s phenomena? It is worth noting that there are several passages in Cretu’s paper that suggest the alternative, antirealist view that “real patterns” are simply defined in terms of empirical data. They are not patterns in the world, but patterns in our data sets. For example: “Real patterns are not authenticated as having a certain nature – that is, as being of a certain kind, having a certain origin, or being constituted in a particular way. Real patterns are authenticated as being genuine phenomena prior to the development of particular perspectives” (2020: 13). The authentication of real patterns “involves nothing over and above identifying, measuring, and maintaining the salience of empirical phenomena” (2020:18). Furthermore, there is a good reason for Cretu to leave this antirealist option open, since when we examine scientific practice, it seems that research traditions often fail to secure ontological commitment.

Consider planetary astronomy. This research tradition authenticated a variety of patterns that have been retained through the Ptolemaic, Copernican, Tychoonian, and modern astronomical perspectives. That is, the empirical patterns that Ptolemy had to account for are still present in the data today, at least approximately so. All of these theories must account for the motions of the Sun, Moon, and planets, as

observed from the Earth. Indeed, part of the argument for the Copernican theory over the Ptolemaic theory was precisely that the Copernican theory better accounted for patterns that were recognized by all, such as the fact that Mercury and Venus are always observed close to the Sun. Yet clearly, given the radically different ontologies postulated by these perspectives, the patterns that are retained are merely patterns in empirical data sets.

In response to this, it might be suggested that the move from Ptolemaic astronomy to modern astronomy constitutes a change of research tradition, not a mere change of perspective. After all, a research tradition includes, as Cretu notes, assumptions about the kinds of things in the world and how they interact. Obviously, our views of the basic entities in the world changed with the Copernican Revolution. But precisely what's in question here is exactly how deep those assumptions go. If the planets are described in more empirical terms, by considering how they appear from the Earth, and the interactions between them that we observe, then the Ptolemaic and Copernican perspectives do share an ontology. The general difficulty here is how exactly we draw the line between research traditions and perspectives. Intuitive judgements about this are likely to reflect what we happen to be interested in. Consider that geocentric models are still used in various contexts: for celestial coordinate systems in positional astronomy, for teaching purposes such as the construction of planetariums, and for navigation. If what matters to us is simply the positions of stars and planets, and we are not working with an extreme level of precision, then the difference between Ptolemaic and Copernican systems looks much more like a matter of perspective. In any case, there is little question that the Ptolemaic and Copernican systems attempted to account for at least some of the same patterns, provided we demarcate patterns with a coarse grain.

A further wrinkle here is that even once a pattern is authenticated, the relevance or significance of the pattern is created by the perspective. Consider again Cretu's example of stellar classification. In 1897, Antonia Maury discovered an anomalous type of blue star. She found that the spectra of some of the brighter blue stars had narrower hydrogen absorption lines, and stronger metallic absorption lines. She designated them with the spectral subdivision "c". At the time, nothing was known about what caused the anomaly in the spectral lines. Maury's discovery was largely

ignored for a decade, until Ejnar Hertzsprung realized that the anomalous spectra were from stars that were more distant, and thus must be more luminous, than those with standard spectra. Maury's pattern revealed the distinction of giant stars. Hertzsprung summarized the impact: "To neglect the c-properties in classifying stellar spectra is nearly the same thing as if the zoologist, who has detected the deciding differences between the whale and the fish, would continue in classifying them together" (quoted in Dick 2013: 96). This led to the Hertzsprung-Russell diagram, which displays the relationship between luminosity and temperature, and is still in use today. The important point for us is that while we take Maury's pattern to track a natural kind, this is not a conclusion based on the pattern itself. Only with the development of a perspective on this pattern – that the narrower hydrogen absorption lines indicated the luminosity class – was the significance of the pattern realized.

I will summarize the problem. Cretu says that real patterns are natural kinds. Real patterns are authenticated in research traditions, and various perspectives can be developed to provide further information about these patterns. Patterns are retained through changes in perspectives. The key idea is that it is the activity of the research tradition that justifies commitment to the real pattern. The problem with this is that there are many cases where this leaves us with a pattern that is far too impoverished to count as a natural kind in any sense, because often what is retained over research traditions are merely patterns in empirical data. Furthermore, this problem generalizes. Suppose there is a research tradition in which a real pattern has been authenticated, and that there has always been a consensus about the underlying nature and constitution of this pattern. The Higgs boson was authenticated in 2012, and there has always been broad agreement about its properties, even before its authentication occurred. Unfortunately, we don't know what perspectives will be developed in the future. In any case, regardless of how perspectives develop, the point is that what the research tradition established is extremely constrained. It is also worth noting that this problem undermines the appeal to indexical redundancy as a way of distinguishing the patterns of science. It only makes sense to talk of measuring a pattern from many points of view if that pattern is treated as something beyond the data. Ultimately then, Cretu fails to satisfy the realist aspect of perspectivism, and fails to show how a perspectivist can justify belief in natural kinds.

Having said this, there is an important insight in Cretu's work. Notice that her account, if it had been successful, would have been compatible with both essence kinds and cluster kinds. It does not actually seem to be a *competitor* to the theories of kinds articulated earlier in this chapter. Of course, we might suppose that essence kind and cluster kind theories are also not competitors to each other, insofar as we could hold that there are both essence kinds and cluster kinds in the world. But suppose we are talking about some specific kind – a species, say. Well then, if it is a natural kind, it is either an essence kind or a cluster kind (or some other type of kind not discussed); it cannot be both. However, we could say both that a species is an essence kind, and that this essence is a real pattern that is authenticated by a research tradition; or that it is a cluster kind, and that this cluster is a real pattern that is authenticated by a research tradition. Hence Cretu's view is not a competitor to the traditional theories. Cretu herself does not seem to recognize this, since she explicitly contrasts her position to essentialism and to the HPC account (2020: 2-3).

As a result, Cretu avoids the Redundancy Problem: her view is doing something different to these earlier theories of kinds. While Cretu's own position fails for the reasons stated above, I think that this point provides a good indication of how perspectivists should proceed with respect to the classification debate. The question, of course, is what exactly we are doing when we construct a philosophical account of natural kinds. Here I slightly modify a distinction drawn by Taylor (2020). Taylor argues that there are two roles a theory of kinds might play. First, there is the taxonomic role, where the theory is used as a kind of "cookie cutter", telling us how the world divides into kinds and non-kinds. Second, there is the metaphysical role, which tells us "what sorts of things natural kinds *are*, how they are structured and what the core features of kinds are" (2020: 2081-2). Taylor suggests that this distinction has been obscured in the previous literature on natural kinds. My own impression is that few philosophers would suppose that *any* philosophical theory could be suited to playing the taxonomic role: as we have seen, philosophical theories are expected to defer to science in these matters. Our best scientific theories provide the "cookie cutter"; then philosophers may propose a metaphysical theory to account for how science tracks the joints of nature.

What a philosopher might well offer, however, is an account of how science executes the taxonomic role. She can ask, what is the nature of the “cookie cutter”? This is how I think Cretu’s account is best interpreted. Here is another way to frame the difference. First there is the question, what do kinds look like, *sub specie* a given theory or research tradition? What is the picture of the world provided by the theory? This concerns the *product* of current science and corresponds to the metaphysical question; we can take this as the question that traditional theories of kinds provide an answer to. We test essentialism, the HPC account, etc., against the results of our best theories. We take it that essentialism fails for biological species because there are, according to our best biological theories, no necessary and sufficient conditions for species membership. Recall the point that we cannot establish Coordination conclusively; we can only “show how coordination could be, by the theory’s own lights and presuming its own truth, something other than a cosmic coincidence” (Franklin-Hall 2015: 932). This is where essentialism, HPC theory, promiscuous realism, etc., play an important role. But we may step back from this, and ask instead, what do kinds look like, *sub specie* the general history of science? How do scientific theories carve up the world? This concerns the practice of science. It is here that perspectivism makes an important contribution. With this point in mind, I turn in the next part of the thesis to the task of developing an alternative perspectivist account of kinds.

IV

PERSPECTIVISM AND KINDS

Chapter 8

Theoretical Kinds

8.1. Introduction

I will begin by summarizing some of the results so far. In Part III, I discussed two arguments for perspectivism, the incompatible models argument and the failure of fit argument. I concluded that there are various plausible realist strategies for dealing with incompatible models, but that failure of fit poses a more serious challenge to traditional realism, and that a number of potential realist responses to this challenge are unconvincing. This leaves open the problem of developing a perspectivist account that respects the insights of failure of fit. In particular, perspectivists need an account of kinds that provides an answer to the Coordination Problem (how is it possible for scientific kinds to be coordinated with the natural kinds?) and the Redundancy Problem (what new contribution does perspectivism make to the natural kinds debate?). Cretu's perspectivism answers the Redundancy Problem, but other difficulties for her account suggest that an alternative is required. Thankfully, an alternative is suggested in Giere's work. In this chapter, I will elaborate on this alternative, and show how it answers the Coordination Problem. In the next chapter, I will return to some of the challenges to the philosophical foundations of perspectivism, in particular Chakravartty's dilemma, and I will show how my perspectivism avoids these problems.

8.2. Giere on theoretical kinds

An intriguing though sadly underdeveloped discussion of a perspectivist view of classification and kinds occurs in a few pages of Giere's book (2006: 84-88), where he introduces the notion of "theoretical kinds". As I understand him, Giere holds that the traditional picture of the role of classification in science is committed to two ideas. First, "the goal of scientific research is to develop concepts that correspond with objective natural kinds" (2006: 84). This is illustrated by my discussion in the

previous chapter: on essentialism, the concepts used in classification schemes should pick out essences; on the HPC kind approach, classification concepts must be accommodated to the causal structures of the world. Even promiscuous realism, which allows a host of cross-cutting classification schemes, still supposes that these schemes are tracking genuine divisions in the world. All these views treat our theories as mapping the world, or as “laying a grid on the world”, as it were: essentialists and HPC theorists think that we will converge on a single grid; promiscuous realists think that a variety of overlapping grids can be laid and we will not converge on a single one.

The second aspect of the traditional picture is an account of theory development that Giere summarizes as “first kinds, then laws, then theories” (2006: 86). We discover kinds in the world, then investigate those kinds to discover generalizations applying to them, and then we systematize and explain such generalizations with more sophisticated theories. I expect that probably no philosopher of science has ever held a view exactly like this, but Giere’s point is primarily just that not enough emphasis had been placed on the role that theories have in determining kinds. Even though it clearly requires theoretical development to distinguish the kinds, these kinds are traditionally treated as among the explananda of scientific theories. For example, there are real divisions in the world that distinguish different species, and the goal of biological theory in this context is to classify these divisions and explain how they came about. Perhaps a more useful way to understand the “first kinds, then laws, then theories” slogan is not as a description of the actual temporal development of science, but rather as summarizing the explanatory and justificatory relations in a mature scientific theory. These two aspects of the traditional picture are then two sides of the same coin: these explanatory and justificatory relations follow from the fact that our classification concepts correspond to kinds in the world.

Giere’s perspectivist alternative reframes questions concerning classification in terms of theoretical kinds rather than natural kinds. What are theoretical kinds? To explain this concept, Giere gives the example of models in classical mechanics. The general principles here are Newton’s laws, and by adding specific conditions we can generate various nested sets of models. The force function $F = -kx$, where x is displacement from equilibrium position, yields the general model for the simple

harmonic oscillator, and then more precise specifications yield yet more specific models: for example, taking x as the displacement of a mass on a spring yields a model of a pendulum. Models of other oscillating systems such as bouncing springs and vibrating strings can be derived by adding different specifications to the simple harmonic oscillator model. Giere then says:

My view is that each of these models represents a *kind* of mechanical system. The equations of motion for these models define kinds of mechanical motion. We therefore have a clear notion of theoretical kinds. They are defined using the principles of the relevant theory. So the theory and the implied laws come first. The kinds are defined relative to the theory. Determining the empirical counterparts of theoretical kinds is another matter altogether. Since no real system ever exactly matches the ideal behaviour of a theoretical model, we cannot simply say that the corresponding set of empirical objects of the defined kind are those whose behaviour matches that of the model. What we count as being empirical members of the corresponding set of real systems depends on how good a match we require and in what respects. These judgements cannot help but be interest relative. (2006: 87)

Broadly speaking then, a theoretical kind is a kind defined in a model. Theoretical kinds are not taken to correspond in any straightforward way with things in the world – one way to put this is to say that theoretical kinds need not (perhaps, sometimes, cannot even in principle) correspond to natural kinds, even on the most liberal view of what constitutes a natural kind. This will be discussed in more detail later. We pick out kinds in the world when we identify elements of the model with elements of a worldly system: for instance, I might identify the bob in the abstract pendulum model with a real mass on the end of a real rod. I can treat the worldly system as if it instantiated the mechanical kind *pendulum*. But for Giere, as we have already seen, models are only ever more or less similar to their target systems, and all models contain idealizations. To say whether the target system instantiates the kind, we must specify the relevant respects and degrees of similarity. In a footnote, Giere adds:

Because theoretical kinds are defined by theoretical principles, and principles are human constructs, my account of theoretical kinds ends up being on the constructivist side of the constructivist/objectivist divide – if one divides up the possible approaches to kinds that way, which I obviously do not. (2006: 131)

Care must be taken in interpreting this. As we have seen, on Giere's view, the general principles are not even treated as representational in themselves, but more as tools for the construction of models. It is models that describe specific entities, and models that are compared to the world. The simple pendulum, the bouncing spring, etc., are not simply defined by general principles – as Giere himself notes, they are not simply deductive consequences of the principles of classical mechanics. Perhaps the better way to express Giere's position is to say that theoretical kinds are defined by idealized models. These models may be more or less similar to systems in the world. But these models also deviate, sometimes very radically, from the real worldly systems. The real system only ever approximates the behaviour of the model in the respects of interest; indeed, it may have radically properties in other respects. Compare how a real gas may approximate the behaviour of an ideal gas with respect to the relation between pressure, temperature, and volume, but is completely different in other respects, such as in terms of its molecular constituents.

It is worth drawing attention to a few aspects of this perspectivist view. First, recall that for Giere, models can be more or less similar to real systems, but no real system will match a model perfectly. Note that it is often the model as a whole, or at least several aspects of the model as a whole, that is compared to the real system. Theoretical kinds are coordinated with the world when the models in which those kinds occur are taken to correspond to a worldly system. So, theoretical kinds are never perfectly exemplified in the world.

Second, even once we coordinate a theoretical kind with entities in the world, there will be much indeterminacy in precisely which entities are members of the kind, because we need to specify what respects of similarity we are interested in and how good of a match is required, and this will depend on the context. Giere gives the example that when specifying mechanical oscillators, if we are running an

experiment with a piece of precision equipment we might require that the period be within one part in a million of the ideal, but if we are designing a child's toy, one part in fifty might be enough (2006: 88).

Third, laws hold exactly in abstract models; theoretical kinds can therefore be said to obey laws (2006: 87). Notice that this preserves the traditional view that kinds are subject to laws, and that identifying something as a particular kind allows us to predict and explain its behaviour in virtue of the laws governing that kind. It is even plausible that essentialism holds in idealized models, that theoretical kinds can be defined as sharing certain necessary and sufficient properties and that these essential properties determine further properties of the kind. It is just that these laws and essences would hold only in idealized models, not in the world. Laws and essences may play an important role in scientific inquiry; where essentialism goes wrong, according to the perspectivist, is in attributing laws and essences to the world.

How does this view play out with respect to the Redundancy Problem? With the possible exception of essentialism, current realist views of kinds already include a significant role for the influence of interests. But what contribution, exactly, do our interests make to our classifications? Well, in these views, interests influence classification insofar as they guide our *selection* of possible classification schemes. There are many kinds in the world, and many cross-cutting classification schemes possible. It would be utterly unwieldy to try to track all the kinds, and so we must decide which ones to latch onto – and this is dependent on our interests since different classifications will be appropriate for different purposes. So, interests play a role in *selection*. It should already be clear that, once we see that science generates theoretical kinds, scientists are not merely selecting from the various kinds and similarities that exist in the world. This is because theoretical kinds are never actually instantiated. Empirical counterparts to theoretical kinds are identified only when we specify: such-and-such features must match, and they must match to such-and-such degree of accuracy. Thus, the constructive element is far more significant.

Perspectives are partially constituted by classification schemes, and the classification scheme is not given by the world itself. As Giere puts it, in response to

the suggestion that social categories are constructed while physical categories carve nature's joints:

If it is correct to understand social categories as defining perspectives within which to interpret social activities, then it is also correct to understand physical categories as defining perspectives within which to interpret the physical world. (Giere 2012)

In the next sections of this chapter, I will examine these constructive practices. I will elaborate on the different types of theoretical kinds, explore their relation to natural kinds, and draw some conclusions for realism and perspectivism about kinds. We will see that theoretical kinds arise both in highly abstract models, but also in the "empirical counterparts" of those models. The reason for this is that theoretical kinds are defined in models, and this includes the data models that are most closely connected to empirical results. Higher-level, abstract models are not compared directly to raw data, but to models of data (Suppes 1962). In science we find hierarchies of models; as Giere himself once put it (1999): "It is models almost all the way up and models almost all the way down." We should similarly expect to find theoretical kinds almost all the way up and down; I will argue that this is just what we do find.

There are at least three classes of theoretical kinds. I do not claim that this taxonomy is exhaustive, and there may not be strict divisions between these classes. This taxonomy is intended to provide a starting point for investigation. Indeed it may be the case that "theoretical kind" – much like "natural kind" – really covers a variety of different concepts. I use the concept "theoretical kind" partly because it provides a useful starting point for an investigation into an aspect of the classification debate that is somewhat overlooked, and partly because, even if "theoretical kind" picks out a host of heterogeneous things, all these things share features that are relevant to perspectivism: in particular, the ways that theoretical kinds relate to the world. Theoretical kinds are interesting because they give content to one of the most controversial aspects of perspectivism, namely the claim that certain properties arise in the interaction of theories and the world. I will now turn to a discussion of the different kinds of theoretical kinds. It is worth noting that in doing so, I may be using

the term “theoretical kind” in a broader sense than Giere, since as I mentioned, Giere’s own discussion of this is brief. Still, all three of these can be identified in Giere’s discussion. With all this said, I will now outline three types of theoretical kinds:

8.3. *Idealizations*

In discussions of objects defined within scientific models, it is standard to draw a distinction between abstraction and idealization (e.g. Godfrey-Smith 2009). In abstraction, certain features of a system are simply ignored within the model. When modelling a pendulum, we may just ignore the colour of the pendulum. In idealization, features of a system are intentionally misrepresented or altered, or the model attributes to the system properties that it does not have. Consider an ecological model of an infinitely large population with non-overlapping generations, or a stellar model where the equation of state is the ideal gas law. Abstraction and idealization are used ubiquitously in the sciences. I will return to abstraction shortly. I mention idealization here primarily to bracket it for the purposes of this discussion. There are two reasons for this.

First, there is already an enormous recent literature on idealization and its connection to problems concerning realism and representation (see for example Odenbaugh 2011, Psillos 2011, Shaffer 2012, Weisberg 2007). Second, more importantly, idealization is not of such immediate relevance to perspectivism. Many idealizations are straightforward falsehoods; in idealizing, scientists are not even attempting to correctly describe anything in the world. Nobody supposes that an ideal gas, a frictionless plane, an infinite population, etc., could be natural kinds, in the sense that they are not candidates for things that might really exist in the actual world. Of course, we can still raise questions about the relation between models containing these idealizations and target systems in the world. But insofar as we identify idealized models with target systems, the theoretical kinds should be understood in terms of the other two types of theoretical kind. The central question for this thesis is whether there are non-perspectival facts. While many people suppose that pendulums exist, that the kind *pendulum* is instantiated by many target systems, even though pendulum models are idealized; nobody supposes that

massless bobs exists. The kind *massless bob* is not instantiated by anything. So even if we should think of the idealization *massless bob* as a theoretical kind, it is not a kind that is relevant to this discussion.

The next two types of theoretical kinds involve processes of abstraction. It is worth noting that I don't think that abstraction, or "abstract objects" in general when these appear in models, should be counted as theoretical kinds. As noted, abstraction often involves simply ignoring properties that are not considered relevant for the purposes at hand, as we might ignore the colour of a pendulum. This is something we do unavoidably all the time, even in non-scientific context – in perception, we never detect all the properties of an object. When we imagine an object, we inevitably focus on some of its properties and ignore others. What is interesting about theoretical kinds is that they involve a constructive process, and this is what we find in the next two types. This is not abstraction-by-omission, but more a matter of abstraction by aggregation of properties (cf. Ordorica 2016 on "abstraction-by-aggregation").

8.4. Schematic objects

These are difficult to define without a detailed example, but essentially, X is a schematic object when X plays some important function in highly abstract models, but no general criteria can be given for what counts as an X. X is coordinated with the world via more specific models that instantiate the abstract model. Consider again Giere's example of a theoretical kind, the abstract harmonic oscillator model which can be used, when specific conditions are applied, to yield more specific models such as the model of a pendulum. These models will involve idealizations in the standard sense: we may assume the pendulum is frictionless, we may treat the bob as a point mass, etc. What is more prominent in this example, I think, is a special kind of abstraction. I take it that what Giere is gesturing at here is that the kind "harmonic oscillator" is coordinated with systems in the world only via the addition of specific conditions to create more specific models; this is his point that Newtonian mechanics provides a hierarchy of models, with the model of the simple harmonic oscillator being highly abstract.

Harmonic oscillators are defined mathematically. In other cases, the schematic objects cannot be defined in this way. I will consider the example of *biological populations*. There are various abstract selective models that involve biological populations. Giere (2006: 71) explicitly states a perspectivist view of natural selection, which is a useful place to begin. Evolution by natural selection occurs in a population when three conditions obtain: (1) There is variation in the traits of members of the population; (2) some traits are heritable, so that offspring resemble their parents more than they resemble unrelated individuals; (3) some of the heritable variation increases the probability of survival and reproduction. This abstract characterization of natural selection comes from Lewontin (1970). Giere describes these as the “principles” of evolutionary theory:

By calling these statements “principles”, I am saying that they are not empirical laws, but that they together define what it is for a population to evolve by natural selection. They define a highly abstract model of natural selection. They do not by themselves say anything about any population of real organisms.

To formulate empirically testable hypotheses, one must designate particular populations and specify which properties of members make them fitter in their common environment. Darwin’s finches provide a canonical example (Lack 1945; Grant 1986; Weiner 1995). The populations are those of several species of finches on the Galapagos Islands. The traits in question are the size and shape of beaks. The environmental features include the available food sources, such as plant seeds. With these specifications, we now have evolutionary models for the evolution of various species of finch. (2006: 71)

In denying that (1)-(3) are laws, Giere is denying that they describe true, exceptionless regularities in the world. What’s really going on here is that these principles provide us with a highly abstract model of a population. If we then consider an actual population, we can use this model to construct a more specific model of our designated population. We need to specify the species of organism, the environment in which the population lives, the particular trait that makes the

organisms better adapted to their environment, etc. This gives us an evolutionary model of that population. So, the principles define a highly abstract model, and then this model is used as the basis to construct new models whose elements are identified with features of real populations. The key point is that what is tested are the specific models, not the initial, abstract model defined by the principles. We cannot say whether the abstract model fits the world. It is the more specific models that may fit or fail to fit the world, for it is only in these specific models that we can identify elements that can be compared with features of real populations.

Something very much like this view of the structure of the theory of natural selection has been developed in greater detail by Robert Brandon (2010). Brandon does not describe his view as “perspectivist”, and his concern is with the concept of adaptedness or fitness rather than with scientific representation, perspectives, truth, etc.; but given the similarities between Giere’s account and Brandon’s, it is worth considering Brandon’s position. Brandon begins by noting that the key explanatory idea in the theory of natural selection is that adaptedness to the environment is what makes an organism more likely to survive and reproduce. When we examine biological populations, we find a correlation between particular traits and reproductive success, leading those traits to become more prominent in the population over time. For example, when we examine a population of birds, we might find that those with longer beaks tend to leave more offspring. What explains this correlation? The answer is that these traits make their bearers better adapted to their environment. It is by making an organism better adapted to its environment that these traits improve reproductive success. With this in mind, Brandon states a foundational law of evolutionary theory, which he calls principle (D). As Brandon notes, (D) is rarely stated explicitly in biology papers or textbooks. The reason for introducing (D) is that the inferences biologists make and the explanations they give in the context of evolutionary theory clearly presuppose something like it:

(D) If *a* is better adapted than *b* in environment *E*, then (probably) *a* will have more (sufficiently similar) offspring than *b* in environment *E* (2010: 107).
(Where *a* and *b* are two individuals taken from some biological population.)

From a perspectivist point of view, (D) defines what we might call the “Darwinian perspective”. It does not in itself represent the world – it does not describe any actual populations. Rather, (D) describes a highly abstract model that is used as a tool to construct more specific models that do represent worldly systems. It is important to see that (D) is not in itself a testable hypothesis. The reason is as follows. Brandon points out that to apply (D), we first need a definition of relative adaptedness; this will tell us what it is for *a* to be better adapted than *b* in some environment. Now, adaptedness cannot simply be defined as actual reproductive success. For one thing, we need to allow that reproductive success might be caused by processes other than selection, such as genetic drift. For another, the theory presupposes that adaptedness *causes* reproductive success; or at least, that adaptedness explains reproductive success. Why did *a* have more offspring than *b*? – because *a* was better adapted than *b*. To identify adaptedness with reproductive success would make (D) a tautology. So we are looking for some property *X* distinct from reproductive success, and which can be found in all biological populations. The problem here is that for any *X* you specify, we can then in principle construct an environment that artificially selects against *X* (Brandon 2010: 114). Since artificial selection is a form of natural selection, (D) is then falsified.

Where have we gone wrong? Either we identify adaptedness with actual reproductive success, in which case (D) becomes a tautology, or we identify adaptedness with some property distinct from actual reproductive success, in which case (D) is easily falsified. There is an alternative: define adaptedness in a way that renders (D) untestable. This is Brandon’s move. He proposes the following definition of adaptedness: “*a* is better adapted than *b* in *E* iff *a* is better able to survive and reproduce in *E* than *b*” (2010: 115). This definition makes adaptedness distinct from actual reproductive success, since the fact that an organism is better able to survive and reproduce doesn’t entail that it will actually have more offspring. But obviously, there is not in general any way of telling whether one organism is better able to survive and reproduce than another. So this definition of adaptedness renders (D) untestable, because we cannot check whether the organisms that are better able to survive and reproduce are in fact the ones that tend to have greater reproductive success. We also, of course, cannot artificially select against those organisms that are better able to survive and reproduce. As we change the selective environment,

we would just change our judgements of which organisms count as better able to survive and reproduce.

(D) is untestable, but it can be used to generate testable models. How do we get specific models from (D)? Well, we add the specific conditions: for “*a*” and “*b*”, we substitute terms referring to conspecifics in a particular population; for “*E*”, we substitute a term referring to an environment; for “better adapted”, we substitute a term referring to a particular trait. For example:

(Di) If moth *a* is darker-winged than moth *b* in forest *E*, then probably moth *a* will have more offspring than moth *b* in that forest.

With these things specified, we have described a representational model, and the elements of the model can be identified with things in the world. In (Di) above, for example, we can identify “moth *a*” with various real moths. This allows the model to be tested. Then what should we do if a specific model (Di) is falsified? Well, we shouldn’t reject (D). Instead, we should simply develop a different specific model, (Dii). As Brandon says, “Consider the instantiation of (D) concerning moths. If through experiments and observations it proved to be false then our response would be and should be that we have incorrectly analysed the ecological situation” (2010: 117). (D) is a highly abstract, schematic model. (Di), (Dii), etc. are more specific models that can be more or less similar, in certain respects and degrees, to real populations. If we want to evaluate (D), the question is not whether it is true or false, but whether it is useful for organizing research.

What does all of this have to do with theoretical kinds? Consider the role that populations play in these selective models. When we instantiate the general schematic model with a specific model, we must assume some sort of partitioning of the biological world so that we have a group of organisms – a population – from which we can draw the individuals we substitute for “*a*” and “*b*”. Having articulated an instance of (D), the question is then just whether this specific model is sufficiently similar to the worldly system. If it is not, then, as Brandon argues, we simply develop a different analysis of the ecological situation. We might try specifying different traits as promoting survival and reproduction, or we might try changing the relevant

environment. Another thing you can try is changing your specification of what counts as the population. We can apply an alternative population concept (cf. Gannett 2003). The abstract model (D) assumes a population, but when turn to the task of coordinating (D) with the world, there is no need to adopt one definition of population. In most cases biologists will just adopt whatever population criterion is intuitive given the species they are studying and their research interests. If what's intuitive doesn't work, try a different classification scheme, and come up with a different model! So, no general criteria for what counts as a biological population can be specified. The boundaries of the population, and how the population is carved up into individuals, are up for grabs.

We have a pluralism about population. There are no general constraints on which population concepts can be used in selective models. But it is a deflationary pluralism, because the biological population turns out not to be a substantive property, but a matter of whatever classification scheme works when generating models that instantiate the schematic model. In the specific model (Di) suggested above, which individuals are candidates for being picked out by "moth *a*"? Well, moths in E. But what are the boundaries of E? What about moths that are sometimes in E and sometimes outside of it? What is interesting about this from the perspectivist point of view is how, in general, the answer to this question simply doesn't matter. We are dealing with simplified, idealized models. There are contexts where successful selective models will draw the boundaries in different ways.

My claim is that "biological population" is a schematic object. The abstract model makes references to a biological population. But to construct the more specific models that represent actual parts of the biosphere – to "designate particular populations and specify which properties of members make them fitter in their common environment" – we must appeal to some specific population concept. There are an indefinite variety of population concepts available, and nothing in the abstract model, or in biological theory in general, can determine exactly which population concept we choose; it depends partly on the actual state of affairs in the ecological situation, partly on our interests and goals, partly on our background knowledge, etc. Crucially, if a specific model is rejected, one possible way to create a superior model is by changing the population concept. The key point is that while "biological

population” is an object used in many abstract models, there are no general criteria for what count as a biological population. Indeed, in contrast to the harmonic oscillator, “biological population”, in general, cannot even be viewed as a mathematical object. It is purely schematic: a kind of empty role to be “filled in” with a variety of potential specifications.

Schematic objects play a significant role in models but have been somewhat overlooked in the modelling literature. It is useful to compare schematic objects with idealizations. Many people in the models literature follow Godfrey-Smith (2009) in treating objects in models as “imagined concreta”: things that would be concrete if they were real. When we reason about ideal gases or frictionless planes, it is plausible to suppose, as Godfrey-Smith puts it, that although they do not exist, “at least many of them might have existed, and if they had, they would have been concrete, physical things, located in space and time and engaging in causal relations” (2009: 101). This is one of the motivations behind the parallel that is commonly drawn to fictional entities. Sherlock Holmes doesn’t exist, but if he did exist he would be a flesh-and-blood person, part of the causal web of the concrete world. The world could have been such that it contained Sherlock Holmes; similarly, it could have been such that it contained an ideal gas.

One of the problems with discussing schematic objects is that we use the same term to refer to both the schematic object and the instantiations of it: see “population” above, which can refer to a schematic population, or to a specific population in a representational model (similarly, we might use “population” to refer to a real population, to a population in a data model, and so on). A schematic object could not be concrete. It plays an important functional role in the model, and it is coordinated with the world via the more precise definition, but we can’t treat it as having a functional role *in the world* because once we have generated many more specific models – once we have instantiated the abstract model – the concept of the schematic object will be different in each specific model. What unites these models is the abstract model that they are instantiating. Schematic objects provide a kind of “filter” through which we see the world. The specific instantiations of a schematic object often do exist: there really are populations, for example; and organisms really are members of one or another population. But we carve out the populations, and we

carve them out differently depending on our interests and depending on what works in the model.

Of course, all of this raises the question of why exactly we need (D) at all. We can test (Di) and judge that it is similar in particular respects to a real population without making any reference to the schematic model. In that case, why not just make do with the various specific models (Di), (Dii), etc.? For the perspectivist, (D) is a perspective on the world: it provides a way of looking at the world that unifies the various specific models. (D) is a tool for constructing specific models, but it's not *merely* a tool. If we only had the specific models, representing changes in particular populations, we would not have a unified theory. Consider how classical mechanics understands phenomena in terms of general concepts such as *force*, *mass*, and *acceleration*. In principle you could model the behaviour of a pendulum without these concepts: the model would just represent the observed behaviour of the pendulum. But then we would have no way of relating the pendulum to other phenomena, such as the motions of the planets. Part of the function of Newtonian principles is to unify the models of various phenomena, by presenting the phenomena as exemplifying certain mechanical concepts and describing the relations between these concepts (cf. Giere 2006: 62). Evolutionary theory, embodied in principles like (D), does the same for models of biological populations.

Various forms of unification play an important role in science. Explanation has sometimes been viewed a kind of unification (Kitcher 1989). More recently, several philosophers have argued that unification as central to understanding. For example, Kosso (2002) suggests that understanding involves grasping the links between ideas, a kind of "conceptual coherence": we achieve understanding when our ideas are connected to form a coherent whole. The representational models in themselves cannot provide this. Compare (Di) with another specific model, say of Darwin's finches:

(Dx) If finch *a* has a longer beak than finch *b* on island E, then probably bird *a* will have more offspring than bird *b* on that island.

What is the relation between (Di) and (Dx)? It is only when we view (Di) and (Dx) as specific instantiations of the schematic model (D) that we have theoretical unification. It allows us to see (Di) and (Dx) as exemplifying the same kind of process. There is a tension between the goals of testability and of providing a general, unifying explanation (Cartwright 1983). Hence the utility of schematic models like (D). As Brandon puts it, “Without (D) there is no theory of evolution, there are only low-level theories about the evolution of certain organisms in certain environments” (2010: 118). By serving as a template for a host of different models, (D) gives us a unified theory.

8.5. Boundary construction

Broadly speaking, boundary construction involves the drawing of discrete boundaries where there is continuous variation in the world. In these cases, it is impossible to satisfy the first criteria for a natural classification scheme noted by Dupre, that of sharp distinctions. Recall Ellis: “we should have to draw a line somewhere if we wished to make a distinction. But if we have to draw a line anywhere, then it becomes *our* distinction, not nature’s” (2001: 20).

8.5.1. Chunking continuous variation

The simplest example of this is provided by geographical boundaries in maps. Consider drawing the boundaries of a wetland (Sismondo and Chrisman 2001), where at no one point is there a clear distinction between the wetland and the surrounding areas, and where the wetland will expand and recede over the year. Sismondo and Chrisman find that in official designations of wetland, “only 8% of the area defined as wetland by one source was defined as wetland by all four” (2001: S45). There can be no precise point at which the wetland becomes the forest, partly because the environment is constantly changing through the seasons, and partly because, even at a single time, the relevant boundary properties are interest-dependent. In distinguishing between the wetland and the forest, we might focus on the gradient in plant species, or we might focus on the gradient of light hitting the ground; different gradients will be relevant to different research purposes (Cadenasso et al 2003: 751). While these two gradients will overlap, they are

unlikely to be exactly congruent. Similarly, Cadenasso et al cite research showing that when demarcating forests, the “forest environmental conditions” often extend well beyond the defined forest, and different boundaries will be drawn depending on which properties are relevant to the given research project. They summarize: “Difficulties in discerning edge effects arise from the fact that each one of several environmental factors may change within a unique edge zone. What may be considered an edge depends on the factor studied, the site specificity of analytic models, and the experimental design” (2003: 755). In general, when demarcating ecological boundaries, Cadenasso et al note that ecologists often appeal to features that may be arbitrary in terms of ecological factors, using items of cultural significance such as fences.

Astronomers face this situation when attempting to distinguish stellar kinds, since the properties that form the basis of stellar classification, such as temperature, luminosity, mass, strength of hydrogen lines, etc., are continuous (Ruphy 2016: 118-124). Earlier we noted the seven standard spectral classes: OBAFGKM, which classifies stars from the hottest to coolest. Each of these is subdivided further into spectral types, using digits from 0 to 9, again for hottest to coolest. Stellar classification is thus dependent on our interests – on exactly how fine-grained we want our classifications to be – and exhibits what Ruphy calls “resolution dependency”: the classification is dependent on the resolution of our instruments, since what looks homogenous at a relatively low resolution will be inhomogenous at a higher resolution. In principle, we could continue to subdivide classes indefinitely, until each class contained just a single star. Such a system would be hopelessly unwieldy. Chunking is unavoidable in cases such as this.

Similarly, there are contexts where astronomers classify stars by mass. The masses of many stars are unknown with any precision, as mass is very difficult to determine; nevertheless, we do know that mass must exhibit continuous variation. Dividing entities by mass is a classic example of an arbitrary classification scheme; Dupre (1993: 17) gives the example of a classification of things into those less than 1kg, those between 1kg-2kg, those between 2kg-3kg, and so on, as one that is intuitively “thoroughly artificial”. But now consider models of stars (Freedman & Kaufmann III 2008: 480): the internal structure of a star depends on its mass. When the mass is

<0.4 solar masses, energy flow is fully convective throughout the body of the star. For stars between 0.4 and 4 solar masses, the star has an inner radiative zone and outer convective zone. For stars >4 solar masses, there is an inner convective zone and outer radiative zone. These differences arise from the fact that the more massive the star, the greater the temperature difference between the inner layer and the outer layer. Note that this cross-cuts the stellar classification given in HR diagrams, which uses spectral type and luminosity class.

Chunking occurs at all levels, including direct observation. Danks (2019: 129-130) cites psychological research on categorical perception. In perception we impose strict categories and lose the ability to detect intermediates. Take /r/ and /l/ sounds. English language speakers cannot distinguish intermediate phonemes – instead, such intermediates are heard as being much closer to either /r/ or /l/. Thus, perception simplifies and idealizes. It presents a simplified picture of the acoustic properties, where continuous variation is experienced as discrete chunks. Then at the very abstract level, chunking is required even when we are explicitly dealing with schematic objects. Consider: “ideal oscillator with bob of more than 1kg”, “ideal oscillator with bob of 1-5kg”, and so on. This would be an arbitrary classification scheme, for the same reason as the same kind of classification applied to real objects is arbitrary.

8.5.2. *Causal/constitution indeterminacy*

When describing some process, we must distinguish between subprocesses that are partly constitutive of a larger process and other processes that are merely causal influences on that larger process. For example, do the gut microbiota count as part of the human organism (what is sometimes called the “holobiont”), or are they separate organisms that interact closely with the human organism? This is the question of under what circumstances two objects count as part of the same biological organism. There are a host of different answers available – a host of different concepts of biological individuality. One answer is genetic homogeneity. On this view, a collection of cells counts as a single individual just in case they are all genetically homogenous. So the *E. coli* cells in the human gut are separate

individuals, because they have completely different genomes to human cells. But there are some serious limits to this criterion:

- Chimerism, where an organism develops from the fusion of two embryos, so is composed of cells with different genotypes.
- The genetic material in the cells of any multicellular organism will diverge over time due to mutation. There will often be errors of DNA replication during cell division. So there is never perfect genetic homogeneity. One way to resolve this problem, and the problem of chimerism, is to require only similarity rather than perfect homogeneity. But then this raises the problem of...
- Monozygotic twins. They have very similar genomes, so would count as two scattered parts of the same organism.
- Moreover, regardless of how strict we take the genetic criterion to be, there is the problem of mitochondria. Mitochondria have their own genetic material, a small circular chromosome separate from nuclear chromosomes. The modern eukaryotic cell probably arose through endosymbiosis; the ancestor of the mitochondrion was engulfed by some earlier cell. So the nuclear DNA and mitochondrial DNA even have separate origins. Yet we do not usually class the mitochondrion as a separate organism.

Of course, the genetic criterion is not the only individuality concept. Numerous such concepts have been proposed in the literature (Lidgard and Nyhart 2017), but they all face similar limitations. There are two important points to notice here. First, *organism* or *individual* are schematic objects in the sense discussed in the previous section. There are many different criteria available, as in the case of *population*. Consider: just as instantiations of (D) require a specification of the relevant population, they also require some means of carving that population into individuals. If some specific hypothesis (Di) fails, we might try proposing a new hypothesis with an alternative specification of what counts as an individual. Second, even having specified a single criterion, we still face causal/constitution indeterminacy. In drawing a line, we are further constructing our theoretical kinds.

One of the fiercest debates arising from causal/constitution indeterminacy is found in the literature on the extended mind. Here is the classic example from Clark and

Chalmers (1998). Inga and Otto both want to visit the Museum of Modern Art. Inga is a normal adult human. She thinks for a moment, remembers that MoMA is located on 53rd Street, and makes her way there. Otto, by contrast, has Alzheimer's, and he always carries around a notebook in which he records important facts. He looks up MoMA in his notebook, sees that it is located on 53rd Street, and makes his way there. Clark and Chalmers suggest that Otto's notebook literally forms part of his memory. It is not simply that when Otto reads his notebook, it causes him to form beliefs about the location of MoMA. Rather, he already held such beliefs (though dispositionally); Otto's beliefs are partly constituted by the scribbles in his notebook. Otto + notebook can be modelled as a single cognitive system.

Notice that the issue here is not that we cannot clearly distinguish between where the brain ends and where the rest of the world begins. Of course, if the scale is small enough, this does become fuzzy, and questions about boundaries will arise, but the point is that in the context of the extended mind debate, this is not treated as a problem. Rather, the question of the extended mind literature concerns which processes count as constitutive of cognition and which are merely causal influences on cognition. One way to put this point is that even once we have distinguished spatial boundaries, or boundaries between different entities, we still face the question of how those entities are grouped into larger systems. Chunking continuous variation is not enough; we must also distinguish causality from constitution.

8.5.3. *Scale selection*

Discrete boundaries can arise as a result of restriction the domain of analysis. There are obvious ways to do this in an *ad hoc* manner. For example, suppose I decide to focus only on the O-type stars and M-types stars. In that case, though there will be continuous variation within each class, there will also be an obvious discontinuity. Of course, this is unlikely to tell us anything particularly interesting. However, there are less artificial ways in which discrete boundaries emerge through our choices. This can happen when we select a particular scale for analysing a given system.

In general, several boundary properties are scale-dependent in that they change with grain size (cf. Strayer et al 2003: 725). Length is the most famous example of this:

the length of a coastline is dependent on the unit of measurement, with larger units of measurement generating apparently shorter coastlines as bends that are smaller than the unit of measurement are “smoothed out”. The length of a coastline cannot be defined independently of a particular measurement system, and the appropriate measurement system will depend on exactly how much resolution we require.

In a similar way, discrete boundaries may be dependent on scale. Consider species. There has long been a debate among biologists and philosophers of biology concerning how to define species, and a variety of different “species concepts” have been proposed. The difficulties here arise from two facts: first, the history of life is one of continuous variation – all organisms evolved from a common ancestor, so there are no discrete boundaries between species; and second, there are a number of different evolutionary factors that can produce a cohesive lineage over time, and there are often just as good reasons to classify based on one factor rather than another (cf. Ereshefsky 1992). Despite early hopes otherwise (such as Ruse 1992), species concepts do not appear to be converging on a single taxonomy. They cross-cut the biological world.

One of the most common ways to distinguish species is the biological species concept (BSC), according to which a species is any population of organisms that can interbreed and that are reproductively isolated from other populations. This concept is clearly useful, and it surely reveals genuine distinctions between populations. Reproductive isolation is a significant evolutionary process. But consider the phenomenon of “ring species”. In a ring species, several populations form a geographic ring: A, B, C, D, E, F; where A and F occupy adjacent areas. Members of A can interbreed with members of B, members of B can interbreed with members of C, and so on, but members of A cannot interbreed with members of F. The standard example is the *Larus* gull, with interbreeding populations forming a ring around the Arctic Circle, terminating in *L. fuscus* and *L. argentatus*, which occupy adjacent areas in Europe but cannot interbreed. Now, in a study that focuses only on Europe in a single time period, the BSC can be very straightforwardly applied in this case. There are clearly defined, reproductively isolated groups of gulls. But “zoom out” to a larger area, and this is no longer an appropriate criterion for demarcating species.

Note also that it would make little sense to apply the BSC across different time periods. I am a member of the same species as humans from 100 years ago, yet obviously I could not interbreed with humans from 100 years ago. We would have to introduce counterfactuals: *if we lived at the same time...* – but then how do we rule out other counterfactuals, which would lead us to judge humans and chimpanzees to be the same species? With respect to technology, it would probably be easier to manipulate things to get humans and chimpanzees to interbreed than it would be to get modern humans and humans from 100 years ago to interbreed. The point is that if we are classifying in terms of the BSC, we already are limiting the range of classification to a particular time period. In all these ways and more, then, biologists can construct discrete boundaries between species, and the conditions for appropriate applications of theoretical species concepts, by altering the “grain size” or the “level of zoom”. In general then, if we take a specific ecosystem at a specific time, there will be clear, objective divisions between the species, with very few intermediary forms. Yet if we “zoom out” and consider a larger area or more of the history of life, objective divisions disappear, and we find only gradual variation. We are forced to adopt a different classification criteria for different purposes.

As I see it, “species” is a schematic object. We have various abstract models that refer to biological species, but to instantiate these models, and construct more specific models that represent some part of the biosphere, we need to plug in some specific species concept. What features of the situation we focus on will influence how we instantiate the abstract model. If we are studying the gulls of Europe, it is possible to plug in the BSC. Not so if we are studying gulls all around the Arctic Circle, where the BSC breaks down due to the failure of transitivity of interfertility. So, boundary construction plays an important role in instantiating abstract models. Coordinating the schematic objects with the world requires further theoretical kinds, in the form of boundary construction.

It is worth noting that there is controversy about the extent to which scale selection is a case of boundary *construction*. Scale-relativity is used in a slightly different way by Ladyman and Ross (2007); and Cretu adopts Ladyman and Ross’s approach in her account of natural kinds. According to Ladyman and Ross, scale-relativity is a feature of the world, so that what exists is relative to the scale at which nature is

measured: “at the quantum scale there are no cats; at scales appropriate for astrophysics there are no mountains; and there are no cross-elasticities of demand in a two-person economy” (Ladyman and Ross 2007: 199). In a later article, Ladyman again says that “there are no water molecules at the scale of quark interactions” (Ladyman 2017: 165). Similarly, Cretu summarizes scale-relativity as follows: “what may exist on a particular spatiotemporal or energetic scale may not exist on a different scale” (Cretu 2020: 3). It follows that the contribution of our practices is limited to what in the previous chapter I called *selection*: we choose the scale, but scales, and the properties that obtain at various scales, are objective features of the world. If this is right, then scale is not really an example of boundary construction. It’s no different from the view that there are many kinds in the world, but we can choose to focus on some but not others – this is very much in line with something like promiscuous realism, for example.

The question is whether it is at all plausible to treat scales as objective features of the world. Is it true that at the astrophysical scale, there are no mountains? That mountains “stop existing” once we zoom out far enough? Or that cats stop existing once we zoom in? Beyond the trivial point that cats are bigger than quarks, it is hard to even make sense of this. It is also hard to square with contemporary science, given that scientific models will not make any distinction, in principle, between scales in this kind of way. Thus, biologists will sometimes appeal to quantum effects to explain biological phenomena: their models will describe causal interactions between quantum and biological (see for example the vibrational theory of olfaction, Hoehn et al 2018). Or consider experiments showing that, in carefully controlled conditions, humans can perceive a single photon (Tinsley et al 2016). It is not that photons do not exist at the scale of human persons, but rather that we are not usually equipped to detect them, and in most contexts, we simply don’t need to. Similarly, it is obvious that the behaviour of a cat will affect quantum phenomena, that astrophysical events may affect mountains, and so on. It seems to me then that scale selection is rather a point about our measurement capacities, our computational abilities, and what is relevant in particular contexts of inquiry. Scales do not exist independently but are constructed by us; and when we change scales, in the sense that we consider features of the world with different properties, we detect different patterns. To

suppose that the scales are independent features of the world is to conflate natural and theoretical kinds.

8.6. *Coordination*

I hope that the preceding discussion demonstrates that perspectivism does have a plausible answer to the Redundancy Problem. What about the Coordination Problem? Well, this problem arises from the assumption that we must take the kinds of science as tracking the natural kinds. Then the question is how we can make sense of this. In developing a theory of kinds, we let science be our guide. Realists about kinds generally accept the view that natural kinds are applied to explain the predictive and explanatory successes of science – as Boyd would put it, our scientific classification schemes and the associated practices are accommodated to the causal structures of the world; this accommodation explains the success of those practices; and that is all there is to natural kinds. This has led to a progressive weakening of the concept of kind, from essences to Chakravartty’s “sociable properties”, as it turns out that the sciences appeal to such a wide variety of similarity relations in their classifications. But it should now be quite clear that the classification schemes found even in our best theories are broader even than this. Science delivers theoretical kinds, and these are applied in induction and explanation; they play a role in “accommodating” our practices to relevant causal structures. But for theoretical kinds, it appears that Coordination cannot be achieved even in principle. Coordination should be abandoned. The Coordination Problem, as originally stated, simply dissolves.

It’s important to note that I am denying Coordination specifically for scientific classification schemes. We might still speak of a theory or model as being coordinated in some sense with the world. Consider again the models of stellar interiors. Stars between 0.4 and 4 solar masses have an inner radiative zone and outer convective zone; stars >4 solar masses have an inner convective zone and outer radiative zone. I can imagine the defender of a view such as Boyd’s HPC theory objecting to my position: “Okay, there is continuous variation in stellar mass so we must decide where to draw the line. In that sense, the classification scheme is dependent on us. But there is a genuine difference *in the world* between different

stellar masses, so we are still attending to objective similarities. We didn't invent mass. Furthermore, our models associate these different masses with various other significant differences, such as the difference in the structure of energy flow in the interior of the star. So we have achieved accommodation to relevant causal structures!" All of this may well be right. However, it does nothing to show how the division we have selected, and how we choose to deal with vague cases right near the boundary, corresponds to anything in the world. I take it that this point is even more clear for other cases of boundary construction and for schematic objects.

Having said this, I will point out that I don't take perspectivism to involve a denial of essentialism, or HPC theory, or promiscuous realism. As noted in the previous chapter, I believe that perspectivism is answering a different question about kinds. In that chapter, I distinguished two questions: (1) What do kinds look like, *sub specie* a given theory or research tradition? That is, what is the picture of the world provided by the theory? (2) What do kinds look like, *sub specie* the general history of science? How do scientific theories carve up the world? With this distinction in mind, it strikes me as perfectly sensible to treat scientific theories as uncovering essences, cluster kinds, and so on. This occurs *within the theory or model*, that is, from the perspective of the theory or model. In adopting the perspective of a theory – in viewing the world from the point of view of the theory – of course we usually assume that the classification provided by the theory carves the joints of nature. That's just what it is to view the world from the perspective of the theory! Thus in a trivial sense, we achieve Coordination.

Consider the chemical elements. The elements might seem to remain a serious difficulty for the constructivist view of kinds outlined here, because if we assume that our best theories are broadly accurate, there really does seem to be a clear fact of the matter what the kinds of elements are, and it looks like essentialism is vindicated in this context. Sorting elements by atomic number is a determinate matter: gold has 79 protons; mercury has 80, and there is nothing in between these, no element with 79.5 protons. Moreover, two elements that differ by only one proton may nevertheless have significantly different properties in other respects. Thus, the elements exhibit sharp distinctions, and there are necessary and sufficient properties for something being a given element. These properties explain the other properties

and behaviour of the elements. It seems then that the chemical elements are clear natural kinds. They are not schematic objects, and there is no boundary construction.

One answer to this is that even if the elements are genuine essence kinds, this does not obviate the need for theoretical kinds elsewhere. As noted previously, the whole point of a naturalistic approach to kinds is to account for the inductive and explanatory successes of science. Even if the “building blocks” of the world are straightforward natural kinds, this does little to account for the variety of kinds used in actual scientific practice (cf. Chakravartty 2007: 161). However, recall that we began the discussion of kinds with two goals. One was to undermine the justification for the realist approach; the second, to develop the perspectivist alternative. The ubiquitous use of theoretical kinds does serve the first goal, at least to some extent, even if natural kinds can be identified in some limited contexts. This is because it would undermine the realist’s claim that tracking natural kinds explains the successes of our best theories. Theoretical kinds are here playing the role of counterexamples to the realist’s success-to-truth inference.

Even so, this would be a significant concession to the realist. After all, it would be open to the realist to defend realism in some other way. Perhaps there is some other justification for the belief that scientific theories track natural kinds, beyond this kind of inference to the best explanation. Moreover, once we have postulated natural kinds, we have a straightforward answer to the failure of fit challenge. Corresponding to the natural kinds is a natural classification scheme; the scheme “fits” particular parts of the world. Then theoretical kinds can be understood as perspectives on the mind-independent world – the world described in terms of the natural classification scheme. We can apply *Escape from Perspective*. This is exactly analogous to how we might apply *Escape from Perspective* to colour perception and understand our colour schemes as perspectives on a world described in terms of light, surface spectral reflectances, and so on. In this case, of course, the perspectivist will say that our theories of light and of surface reflectances must also be understood perspectivally (cf. Giere 2006: 40). But now the debate just shifts to how we should interpret our theories of colour science. If at some point we apply *Escape from Perspective*, and up at descriptions that are not interpreted perspectivally, then

perspectivism is undermined. This is of course why the failure of fit argument is supposed to apply to all domains.

In my view, the chemical elements are essence kinds, but they remain theoretical kinds, not natural kinds. First, scale construction plays an important role in distinguishing the elements as we do. Bursten (2018) makes this point in her discussion of the properties of gold. What are the properties of gold? Suppose we have several billion billion atoms of gold. If these are grouped into one lump, then we will have a standard piece of gold: a yellow, shiny, solid lump, that can be worked into different shapes, that is resistant to corrosion, and so on. On the other hand, suppose we have many clusters containing a few hundred atoms each, suspended in solution to prevent them from coagulating back into a single lump. The physical properties of each cluster, such as ductility and conductivity, will differ significantly from the properties of the large lump. These properties will change, as we change the size and shape of the clusters. Finally, if we have a collection of individual, non-interacting atoms in a vacuum, they will not have any macroscopic properties at all. The microstructure underdetermines the macroscopic properties.

Bursten then states a dilemma. On the one hand, we might treat all these collections of atoms as instantiating the kind *gold*. In that case, knowing that something is gold does not tell us about its macroscopic properties. This seems like the wrong result, not least because it is precisely knowledge of the macroscopic properties of gold that allowed us to make inferences about its underlying microstructure. More generally, when distinguishing some kind *K*, we take it that members of *K* share various properties. If *K* is an essence kind, the essence should play some role in explaining these properties. But on this horn of the dilemma, where we are treating these different collections as instantiating the same kind, it turns out that there is only the essence, and no further properties to explain. The most that can be done is to describe the microstructure.

On the other hand, we might treat these collections as instantiating different kinds. This is in line with scientific practice: as Weisberg (2006) notes, in practice, chemists individuate kinds by appeal to structure and reactivity on at least three levels: atomic, molecular, and molar. In which case, it is quite natural to talk about different kinds of

gold. But on this approach, the microstructure is no longer playing the role of an essential property that distinguishes one kind from another. We distinguish essences only by constructing a scale on which we treat elements as individual atoms.

This point can also be seen when we consider how we apply our classifications of elements at other scales. It is a trivial point that, in practice, we never have a perfect sample of any given element, nor can the properties of any sample be precisely specified. Although trivial, this has significant consequences for our thinking about natural kinds, as elaborated by Chang (2004) and Teller (2018b; 2019). Take the first case described earlier, where we have a large lump of gold. Suppose we want to determine the properties of such materials. We might ask, what is the melting point of gold? Well, 1064°C. But obviously, any actual lump of gold will have some impurities. So here we must decide what counts as a “pure” sample. Moreover, properties such as melting point will depend on environmental conditions. We might fix all of this precisely, so that we say something like: the melting point of gold is X, where gold is defined as _____, and where conditions A, B, C... etc., obtain. Now, however, this will tell us almost nothing about the actual world, since “gold” in this sense almost never actually exists. The kind “gold” turns out never to be instantiated. It is a theoretical kind. There is no one property that can be specified as *the* melting point of gold. To have an invariant melting point, you would need conditions that do not obtain: a perfectly pure sample, of a particular size, with the atoms grouped together in a particular arrangement, and in fixed environmental conditions.

The elements, as classified in the periodic table, are one of the best cases for a realist account of kinds. But what allows the elements to exhibit the three desiderata of a realist approach to classification is that we have taken a particular perspective on the world. That is, the kinds of the periodic table come into view only by constructing a particular scale of analysis. We must decide to treat elements as single, non-interacting atoms. This tells us nothing about the ways in which elements are distinguished in other contexts – in other sciences, in engineering, in everyday life. Recall how zooming in on an ecosystem reveals clear discontinuities between species. Along the same lines, zooming in on the “building blocks” reveals clear discontinuities. What we cannot do is take the further step of supposing that this scale provides the “correct” perspective, or is the “ideal” perspective, or anything

along those lines. There cannot be a correct perspective with respect to classification.

This point is helpfully illustrated by an example used by Van Fraassen (2008) in a somewhat different context. Van Fraassen gives the case of Professor Deerstalker, who is presenting a theory of the factors affecting deer population in an area of Princeton. Deerstalker's data model is a graph of the growth of the deer population. Deerstalker claims that theory T provides models that fit with the graph presented. Now, it appears that a challenger can grant that the theory fits the graph – that is, that it fits the data model – but still ask whether the theory fits the actual growth of the deer population. The challenger's question is: does the graph provide a good model of the actual growth rate of the deer population? This might be taken as a request for clarification on how the data was acquired. But this misses the point. The challenger can grant that the measurements were made diligently, and that the theory fits the outcomes of the measurements, but still question whether it fits the actual deer growth. The point is that the real deer population growth is distinct from a graph, distinct from a data model. It is distinct from any representation, and a challenger can always ask whether any given representation is a correct representation of the real thing.

However, van Fraassen says that Deerstalker *himself* cannot entertain this doubt without withdrawing his graph. To grant that the theory fits the graph, while doubting that it fits the actual growth rate, but still maintaining the graph, would be analogous to saying: "the deer population growth is X, but the sentence "the deer population growth is X" is not true, for all I know" (cf. van Fraassen 2008: 256). For van Fraassen, there is, for a given individual proposing a representation, no distinction between the population growth and the representation of the population growth. However, this does not change the fact that multiple perspectives are possible. Somebody else might present a different graph, perhaps even a different graph that is equally well supported.

From a third-person point of view, it is true that to say that a theory is adequate to the phenomena is not the same as saying that it is adequate to the phenomena as represented by Deerstalker. That is, to say the theory is adequate to P is not the

same as saying it is adequate to Deerstalker's-representation-of-P. But Deerstalker himself cannot make this distinction in the same way. The reason is that in checking any claim of adequacy, we are always comparing one representation against another. That is, there is always some data model against which the theory is checked; the theory is never confronted with "raw phenomena". The most we can do is check the theory against our representation of the phenomena. Thus, van Fraassen concludes: "*For us* the claim (A) that the theory is adequate to the phenomena and the claim (B) that it is adequate to the phenomena as represented, i.e. *as represented by us*, are indeed the same!" (2008: 259). This is as a matter of pragmatic tautology. To put this another way, suppose I offer a data model D of some phenomenon P, and I take D to be an adequate representation of P. Another person can object that, while D is embeddable in the theory, it is questionable whether P really conforms to the theory. They can claim that there is distinction between D and P. But *I* cannot do this without involving myself in a pragmatic contradiction.²

The problem with van Fraassen's position is that given the context, there is presumably a fixed answer as to how many deer there are. Consider Deerstalker, confronted with the question of whether his graph of the deer population is accurate. Deerstalker can imagine "ideal circumstances" for counting deer. He can imaginatively project a perfectly reliable deer-counter onto each patch of ground at a particular time, over a long period, and imagine that the graph was constructed by such counters. This is an "ideal perspective" for counting deer. This perspective itself rests on various assumptions – for example, we must decide what counts as an individual deer. Nevertheless, with enough such assumptions in place, this perfect perspective will deliver an answer about what the real population growth is. Then we

² Van Fraassen treats representation a morphism between model and target system. But for reasons discussed in the context of the mathematization response to the failure of fit argument, the phenomena do not themselves enter the morphism, since the phenomena are not in themselves mathematical objects. We therefore need to construct data models. Now van Fraassen's move is to simply accept this point, but then argue that, for a particular individual in a particular context, there is no difference in the representation of the data and the representation of the phenomena.

can compare this to our graph, which is, of course, constructed from a far less-than-perfect perspective. Deerstalker recognizes his own cognitive limitations. So he can say that the theory fits the graph, and say that the graph meets the standards required for this kind of data, while still insisting that the theory might not fit the actual population growth.

Van Fraassen ends up in an awkward position in this case. Clearly Deerstalker's estimates could be wrong, and shouldn't Deerstalker himself admit this? Uncertainty about data models is often expressed by scientists. Contessa (2010: 518) states the dilemma: either admit there is a fact of the matter about the number of deer, independently of the data model; or deny this, and treat number of deer as determined by the data model. By taking the first horn, we introduce a genuine distinction between asking how well the theory fits the data, and how well the theory fits the world, even from the point of view of the person offering the data model. By taking the second horn, we embrace metaphysical idealism.

In the example of the graph of deer population growth, we can imagine an "ideal perspective" for the construction of the graph. That is, we can imagine somebody counting each and every deer. So we can doubt that the data model we have actually constructed, on the basis of far less extensive measurements, is in fact accurate. What we are imagining here is somebody who shares our perspective – who shares our classification scheme, our way of carving up the world – but who has far better capacities and resources. The ideal perspective is an ideal form of our own perspective. But now consider the case of boundary drawing. Take a simple example such as drawing a line between two areas. When studying the deer population, we need to specify: what are the boundaries of the environment from which the deer population is drawn? How do we count a deer that is sometimes within these boundaries, and sometimes outside them? And so on. Again, assuming there is agreement on these matters, we can project into an ideal perspective and ask whether we have correctly counted all the deer. But we cannot project into an ideal perspective on boundary drawing and ask whether we have correctly carved up the relevant kinds. There is no ideal perspective to be had on this – at least, not without giving up the classification system altogether. The "ideal perspective" on the biosphere gives up biological population and just sees a series of individuals...

except, of course, it doesn't even see that, because we must construct the boundaries between individuals, as already discussed. There is no ideal perspective available here.

Similarly, consider drawing boundaries between different types of star. Suppose I have a model of the distribution of stars in the Milky Way galaxy, showing how the proportion of particular types changes with distance from the galactic centre. I present a theory of the factors affecting the distribution of spectral types throughout the Milky Way. I claim that this theory provides models that fit with my data model of this distribution. Now comes the challenger, who grants that the theory fits the data model of the distribution of spectral types, but questions whether the theory fits the actual distribution of spectral types. After all, the actual distribution of spectral types is distinct from a data model. But notice that what is shared by both the challenger and myself is the classification of stars into spectral types. While it makes sense for me to suppose that I may have miscalculated the proportion of, say, G-type stars in a particular area of the galaxy, or that I may have misclassified an F-type star as a G-type star, it is less clear how I would doubt the Morgan-Keenan classification scheme. What would it mean to talk of "error" here? Some classification schemes are, of course, more or less useful than others, depending on one's purposes. But no organization of entities can in itself be incorrect. Again, I can imagine an ideal observer in ideal circumstances giving a different count of the G-type stars. But to attempt to adopt an ideal perspective on the stellar classification scheme would be to give up stellar classification entirely, and just take all stars as unique individuals on a spectrum. As with the deer, this would not be an ideal perspective in general, because it would still be up to us to specify what counts as a star, where the line is between a star and space, and so on.

Contra van Fraassen, I think it is coherent for Deerstalker to hold that his theory is adequate to the phenomena as represented in his graph, while doubting that it is adequate to the actual phenomena. Perhaps the deer have been miscounted. Nevertheless, van Fraassen is onto something; certain aspects of our representations are not open to doubt in the same kind of way. This is because when we try to project into an ideal perspective with respect to the classification schemes themselves, we simply lose the ability to classify at all. To adjust van Fraassen's

quote: “*For us* the claim (A) that the theory is adequate to the phenomena and the claim (B) that it is adequate to the phenomena as classified, i.e. *as classified by us*, are indeed the same!”

8.7. Social constructivism

Before closing this chapter, it is worth making a few comments on social constructivism. It may be objected that the account of kinds presented here entails a commitment to a problematic form of constructivism. Clearly, this is constructivist in some sense. This constructivist element was already acknowledged by Giere, though Giere is skeptical that the distinction is useful: “my account of kinds end up being on the constructivist side of the constructivist/objectivist divide – if one divides up the possible approaches to kinds that way, which I obviously do not” (2006: 131). But if kinds are constructed by us, and if we take the plausible view that the world consists in a collection of various kinds of thing, then in what sense do we have access to an objective world at all?

Although classification schemes are in an important sense socially constructed, they exhibit an important feature that Hudson (1994) calls “volitional impenetrability”.

Consider these propositions:

- The Sun is a G-type star.
- G-type stars live approximately 10 billion years in the main sequence.

Are these statements true? Yes, obviously. What makes these statements true?

Well, they are made true because of the way the world is. Having adopted a particular classification scheme, based on particular theoretical principles, it is then largely out of the hands of scientists where exactly the chips will fall. It is logically possible that scientists could take the Sun to be an O-type star instead, but doing this would require radical theoretical changes. But even this has no special connection to perspectivist account of kinds; rather it is the familiar point about underdetermination: we can defend pretty much any proposition we want, provided we are willing to alter enough of our other beliefs. The important point is that not just any proposition can be accepted as part of a consistent and predictively powerful theory that is anything like our currently accepted theory, and this remains the case on a perspectivist view.

Nobody can simply decide to classify the Sun as something other than a G-type star. Even if everybody believed that the Sun were an O-type star, we would still be wrong; it would still be demonstrable that, within the scheme we are using, the Sun is a G-type star. So there are two points that mitigate concerns about constructivism. First, once a classification scheme has been chosen, and once a population of entities has been carved out, we have very little control over which class those entities are sorted into. Given the population of stars, the Morgan-Keenan classification of that population, there is no question that the Sun is a G-type star. From within the perspective of a given scheme, we have control only at the boundary cases. This is one aspect of volitional impenetrability. Second, we will find that many classification schemes simply do not suit whatever our purposes are. Take the instantiation of schematic objects. Here there is conceptual pluralism: the schematic object “population” can be instantiated by many different population concepts, to generate the more specific representational models. But we will find that in particular contexts, it would be extraordinarily difficult to select anything other than the population concept that we do in fact select. Or consider the case of species. Imagine trying to apply the BSC to the entire ring species of Arctic gulls. Even if we find a way to shoehorn the populations into that system, there are alternatives that better capture the evolutionary relationships. In other cases, there is just no way that this criterion can get any purchase, as with asexual organisms.

There are many ways in which classification schemes can fail to fulfil the roles they were constructed to serve. Sometimes, this due to limitations on our part. Taylor (2021) cites the example of phenetics in biological taxonomy. Phenetics sorts organisms on the basis of similarity relations. The aim of phenetics was to construct groupings that are theoretically neutral, that do not presuppose any particular theory about the processes of evolution, so that these could serve as unbiased data by which alternative evolutionary theories could be judged. This ran into the problem that judging similarity between organisms requires us to specify particular properties and to weight them. But which properties we take to be relevant, and how they are weighted, is always driven by background theory. Indeed, there is no alternative, since there are an infinite number of ways in which two entities might be similar or dissimilar (cf. Goodman 1972), but we cannot enumerate an infinite list. We need to

draw on background theories to select the relevant properties. Phenetics could not fulfil the role it was designed for, so alternative classification schemes were developed. Clearly then, it is not the case that we can simply choose perspectives, in the way that we might choose what colour shirt to wear.

Hacking (1999: 6) suggests that the central claim of social constructivism is that to say that some X is constructed is to say that “X is not determined by the nature of things; it is not inevitable.” Again, take stars. If we consider the stellar classification scheme as a whole, Hacking’s slogan this is certainly true. We need not have drawn the lines where we did. We need not have focused on the properties that we did. Furthermore, we may grant that there are various properties that we were “pushed” by our environment and our cognitive capacities to take note of – for instance, we can determine the intrinsic luminosity of a star much more easily and more reliably than we can determine its mass, so it’s unsurprising that the former is more important in our stellar taxonomies. This is a result of our place in the world and our abilities; it is conceptually possible that there could be beings who can more reliably determine mass rather than luminosity, and their classifications schemes would likely thereby be different. In this respect, stellar kinds are not inevitable.

On the other hand, given a particular classification system, and a particular research context, there is often very little doubt about how the world is to be pictured. This is volitional impenetrability. The Sun is inevitably a G-type star. Importantly, volitional impenetrability does not hold for social objects such as money. First, we can simply decide that a particular type of coin is no longer worth £1. This happens when the currency is changed. Perhaps a particular coin is too easy to counterfeit, and is then updated. The old £1 coins are now worthless; they are just lumps of metal. This cannot be done for stars, at least not without displacing the classification scheme. Second, what £1 is changes, in ways that are dependent on how people act: for example, the value of the coin changes due to inflation. This is not the case for a particular star. Its journey through the constructed categories – G-type to red giant to white dwarf – is independent of us.

Chapter 9

Escape from Perspective

9.1. Introduction

As we have seen, perspectivism is a rich family of views; the term is used in different ways by different philosophers. This leads to difficulties with evaluating perspectivism in general. My approach in this thesis has been to focus specifically on the form of perspectivism proposed by Ron Giere, and more precisely, on the view that science delivers knowledge only of perspectival facts, where a perspectival fact “is a proposition that is only true from, or within, or relative to, a given perspective (or limited set thereof)” (Chakravartty 2010). In the previous chapters, I have presented the case for taking perspectivism in this sense seriously. In this chapter, I will examine in detail the central challenge that has cropped up regularly throughout the thesis: the threat of Escape from Perspective. I begin with a brief review of the argument from the previous chapters.

What can it mean to say that truth is relative to perspective? The guiding intuition of perspectivism is the notion that truths are revealed from multiple viewpoints. This intuition may be elaborated in various ways; the standard motivation, formulated in the incompatible models argument, was found to be unconvincing because there are numerous ways that realists can account for incompatible models. The more powerful argument is the argument from failure of fit. I have tried to show that this presents a genuine challenge to traditional realist approaches. I have argued that the realist responses to failure of fit are unsuccessful, and in the course of critiquing the primary realist response, the appeal to natural kinds, I elaborated a perspectivist alternative based on the notion of theoretical kinds. It is important to be clear about the scope of the failure of fit argument. Failure of fit is a challenge: if a realist account of kinds fails to adequately meet this challenge, this leaves us with only a draw between realism and perspectivism. The perspectivist needs some positive reason to suppose that there is a failure of fit in the relevant sense. This is provided by the perspectivist account of kinds.

To give a summary: Facts are true propositions. Any proposition about the world will assume some classification scheme, some way of carving up the world. Such schemes are provided by perspectives. That is, part of what constitutes a perspective is the means of classification. There is no perspective-independent standpoint from which such schemes can be evaluated, since evaluation always takes place with some scheme in mind. Furthermore, and crucially, these schemes do not reflect nature's joints. Human interests are involved throughout the construction of theoretical kinds: interests do not simply guide the selection of admissible schemes, as the realist would have it. Then given a particular framework for carving up the world, various regularities can be identified and true claims can be made.

As was noted in the introductory material, perspectivism is in some ways tangential to the standard realism debates. Perspectivism is not a form of limited realism analogous to structural realism, entity realism, or constructive empiricism. The perspectivist is silent on which propositions are true and focuses instead on the question of what it is for our representational tools to relate to the world such that claims generated on the basis of those tools can be said to be true or false. The perspectivist tries to understand the representation-world relation in a different way to the standard realist. One can be a perspectival structural realist, a perspectival entity realist, a perspectival constructive empiricist. So, we took it that a condition of adequacy for a perspectivist theory is that it needs to be compatible with these different approaches to realism. The view presented in the previous chapters achieves this. For instance, the perspectival empiricist might hold that science delivers knowledge only of the observable world, and that the facts about the observable world are perspectival facts. I will now turn to the problem of Escape from Perspective.

9.2. Escape from Perspective

Looking at the objects in front of me, I can say that my computer is to the left of my phone. This is true from my visual perspective. In some sense, this is a perspectival fact. From a different perspective, my computer is to the right of my phone. Yet there is no interesting kind of perspectivism here, because we can step outside the

confines of specific visual perspectives, and understand how the spatial relations between the computer, the phone, and the observer ground these facts. This is Chakravartty's (2010) challenge, which I have called the challenge of Escape from Perspective. Or at least, that is the intuition behind Escape from Perspective, but there are different ways to elaborate this argument, which pose different challenges depending on the form of perspectivism we are considering. In particular, I think there are three related anti-perspectivist arguments here.

Of course, the failure of fit argument is in itself designed to challenge the Escape from Perspective move. What is unique about failure of fit is that, at least in the strong form presented by Kuhn and Giere, it generalizes to all propositions, whether they are scientific or non-scientific, whether they are theoretical or observational. In that case, by "stepping outside" of one perspective, we will inevitably find ourselves in a domain that can itself only be understood perspectively. Either truth and facts are constituted within perspectives – where this is simply the nature of truth across the board, including for truths of our best-confirmed theories – or we face radical skepticism, a total suspension of judgment concerning any empirical facts. This is the general response to the Escape from Perspectivism intuition. But it is, I think, still worth examining the arguments in some detail.

9.2.1. Distortion again

First, there is Escape from Perspective via the specification of distortion. As perspectivists often point out, all our models are idealized, partial, and simplified. In these respects, our models introduce distortions. But this is no challenge to standard realism when the context and the broader theory allow us to specify the nature of the distortion. Recall the discussion of instruments: it doesn't matter how the instrument transforms the input, provided we have an accurate model of the transforming processes. No matter what the output of an instrument is, we will only take it as "distorted" if our model of the operation of the instrument is inaccurate. Of course, certain instruments may be more or less easy to model, or may produce more or less intuitive representations of worldly systems, but these are pragmatic matters. In the same way, we can use distorting models, without any threat to realism, whenever we understand the respects of the distortion. This is the general idea of Escape from

Perspective: we can step outside a given model and consider how the model relates to the world from a broader model or theory, and this will allow us to identify aspects of the model that are accurate.

Indeed, our very ability to talk about idealizations in the first place demonstrates that we can make such discriminations. How could we know that some aspect of a model is idealized, if we did not have some means of specifying the ways in which the model distorts? Consider a stellar model that treat a star as composed of an ideal gas. We take some parts of the model as accurate, such as its representation of the structure of the convective and radiative zones of the star. We take some parts as distorted, such as its description of the behaviour of the molecules constituting the star. Now clearly, the model *in itself* makes no assumptions about which aspects are accurate and which are idealized. Had science developed differently, we might have assumed that stars really were composed of ideal gases, and in that case, this aspect of the model would not be treated as an idealization. All the work here is being done by our interpretation of the model, by our broader understanding of how the model relates to its target system. So distortion is not a problem: all that matters, the realist will say, is that we have the right interpretation of the model. If we don't have the right interpretation, then either we are just mistaken (in the case of an incorrect interpretation), or we remain agnostic (in the case where we don't make any assumption about what the right interpretation is).

In his discussion of the interpretational approach to representation, Nguyen (2020: 1025) gives the example of pictures of Obama, one with its colours inverted. The original picture is more similar to the real Obama in terms of visual appearance. But which is the more accurate? Neither. Indeed, they have the same representational content. For the original picture, we use an interpretation function that utilizes similarity relations to generate claims about Obama: the fact that a certain area of the picture is light and another area dark allows us to infer that Obama was wearing a light shirt and dark jacket. We can generate the same claim from the inverted picture, *using a different interpretation function*. Nguyen summarizes:

We can reconsider the presumption that idealized models are misrepresentations in the first place. Rather, we should direct our attention to

the claims generated by such models in scientific practice. If those claims are true, then the models themselves can be justly considered accurate representations, despite being idealized.

Similarly, consider maps again. It is tempting to say that the Mercator map represents Greenland as bigger than Australia. But this is a mistake. We know how the map was constructed, and we know how to read it, so we know what it does and does not say. It tells us many things about the relationship between Greenland and Australia, but not that the former is larger than the latter. Indeed, it expresses exactly the opposite. Giere himself notes the interest relativity of maps (2006: 73), and the goals of the users and creators of the map partially define what maps express. Contra Giere's later statement that "these projections are all incompatible", when read properly, there is no incompatibility. They say much the same things. Why then choose different maps? Well, there may be aesthetic reasons, or political reasons; of more immediate practical importance, maps that look different may be associated with particular activities, as the Mercator's use in navigation. Two maps may express the same propositions, yet some propositions may be more salient than others.

So, one way to apply Escape from Perspective is to shift to a broader theory or model that provides an interpretation function specifying the nature of the distortion in some other theory or model. There are important differences in how these interpretation functions may be applied to any given representation. A simple harmonic oscillator model represents a system as a point mass or as free from friction. The Mercator map does not represent Greenland as being bigger than Australia, even though area labelled "Greenland" is bigger than the area labelled "Australia" on the map. But in both cases, we have Escape from Perspective, because we understand the representational norms and we understand how to use these representations to generate accurate claims about the target systems. Now, it should be clear enough that this approach is not generally applicable to the scheme-relativity approach to perspectivism. When we specify the respects in which a model distorts, all the work is being done by the interpretation function. But this requires an appeal to broader some theory or framework, to which the more general failure of fit argument applies. That is, if we adopt a perspectivist analysis of this theory, then the fact that this theory specifies how some model relates to a target system will not in

any significant sense amount to escaping perspective. We may have good grounds for supposing that the theory is true, and we may have good grounds for supposing that on the basis of this theory we can reliably distinguish the accurate and idealized aspects of the model. But truth here need not be understood in a traditional realist sense.

9.2.2. Dispositions

The second approach to Escape from Perspective is found in Chakravartty (2010), who appeals to dispositions. Chakravartty argues that dispositional realism undermines perspectivism. According to dispositional realism, we have knowledge of the underlying causal properties of systems, and these properties manifest differently in different conditions. As Chakravartty puts it: “what one learns via scientific investigation into target systems in the world are their modal features: how having certain properties dispose them to be or to behave in different contexts” (2010: 409). We have already seen this with instruments. Our models of instruments describe how the same causal process, e.g. decay of aluminium-26, can be measured in different ways, thus producing different representations of that process. In the chapter on instrumental perspectivism, I argued that instruments, among various other things in the world, are well-ordered causal processes, and different causal processes will naturally provide different “perspectives” on phenomena involved in those processes. Once we grant that the world is a causal web, dispositional realism becomes very attractive.

Dispositions are intrinsic properties of systems, not relational. They are properties that “dispose the systems that have them to behave in particular ways in specific circumstances.” Chakravartty uses the example of salt (2010: 409). Salt is soluble in water. But salt is not always soluble in water: if the water is already saturated, or if it is in a strong electromagnetic field, the salt will not dissolve. The dispositional realist accounts for this by saying that salt has a property that disposes it to manifest solubility in X-circumstances but not Y-circumstances. What is more, we are now capable of describing this property in some detail, in terms of models of the molecular structure of salt, inter- and intra-molecular forces, and so on. Recall that a perspectival fact is a proposition that is true only within a perspective. The

dispositional realist argues that there are no perspectival facts, only dispositional facts: dispositions that cause systems to behave differently, in different conditions. Alternative descriptions of the same target system can both be true, because manifestations of the same property can vary depending on the circumstances. In a sense, we must always take a perspective on a target system in order to investigate it. But how the target system interacts with or manifests itself within that perspective is not a perspectival fact. Our best models give us non-perspectival information about the dispositions of target systems.

Stated like this, there is a swift answer to the dispositionalist response.³ We can grant that the dispositionalist raises a serious challenge to instrumental perspectivism. The output of an instrument is dependent on the internal structure of the instrument, but there is no reason to suppose that this prevents the instrument from providing a reliable representation of the properties of the target system. It all depends, as we have seen, on our instrumental model. What makes perspectivism apparently appealing in the case of instruments is that target systems will interact with different instruments in different ways, but we can explain this by appealing to dispositional properties of the target system. Turning to theoretical perspectivism, and the argument from failure of fit, we are no longer dealing with interactions between two physical systems. In this case, there will be situations in which the same target system, in the same conditions, is represented in different ways. It doesn't seem that we could account for this by appealing to dispositional properties of the target system, firstly because these different representations result partly from our interests in representing the system, and secondly because the dispositional properties themselves will be represented differently for the same reason.

There is, however, a natural extension to the dispositionalist argument, which is at least suggested by Chakravartty (2010: 410), and this is to think of dispositions as

³ Massimi (2012) gives a more detailed response to dispositional realism. She argues that dispositional realism faces a problem of epistemic bootstrapping that leaves it unable to account for how our knowledge of dispositional properties is justified. As I have discussed previously, Massimi's perspectivism focuses on the epistemic question of how scientific knowledge is developed. She raises no objection to the dispositionalist's denial of perspectival facts.

supporting different manifestations under different conceptualizations. Or put differently: the same dispositional property may manifest differently under different perspectives. That is, property P falls under class C according to perspective A, falls under class C* according to perspective A*, etc. Two points can be made against any general perspectivism. First, the facts about how a given property manifests within a perspective are non-perspectival, since they will hold regardless of what perspective we take. Both the Newtonian and the Einsteinian can agree about how a target system looks from a Newtonian perspective, or how a target system is represented within the Newtonian perspective. Indeed, if they could not agree, at least broadly speaking, on this point, then there would be no way for them to even understand the opposing theory. I doubt that many perspectivists would want to be committed to such an extreme incomparability thesis: it is unlikely that even Kuhn endorsed this position (cf. McMullin 1993) and it is explicitly rejected by Giere (2006: 82-84). So it seems that there are non-perspectival facts about how the world manifests within perspectives. The same point is made by Votsis with respect to colour categorization:

It is of course true that from the perspective of a normal human trichromat the sky appears blue. But *this* truth does not change when considered from the perspective of everyone else, i.e. normal non-humans and normal human non-trichromats, who is cognitively competent with respect to the given task. (Votsis 2012: 92).

Second, just as we can describe the underlying categorical facts that dispose salt to behave in particular ways under different conditions, so we can describe the facts that dispose target systems to manifest in particular ways under alternative perspectives. Humans are part of the causal web of the world, and we can model humans as being very complex measuring instruments. The perspective adopted by a person, including the classification scheme they use to interpret experience, and the “output” in the form of their statements about the world, are all part of the instrument. Consider two societies with different colour classifications. Society S distinguishes colours in the same way that modern English speakers do. Society S* distinguishes only four basic colours: white, black, “warm” (red, orange, yellow), and “cool” (blue, green). We can describe the underlying facts about surface reflectances

and transmission of light that, in conjunction with the human visual system and a particular cognitive framework for interpreting visual perceptions, will dispose humans to utter statements such as “the sky is blue” or “the sky is cool”. So from a dispositionalist point of view, we have non-perspectival facts in the form of propositions about how target systems manifest under different perspectives, and also propositions about the properties of target systems that lead to these different manifestations.

The question is why we should accept that there is any non-perspectival fact concerning how some target system appears from (or manifests within, etc.) a given perspective. To simply assert this as an objection to the perspectivist would be plainly question-begging. Giere is explicit that perspectivism must be applied reflexively to itself (2006: 95). What exactly is the “Einsteinian perspective”, for example? Any classification of perspectives, or other cognitive constructions such as models, theories, and so on, will itself be subject to the arguments of the previous chapters. Indeed, notice that the dispositionalist, in order to make her objection to perspectivism, needs to take a particular view of humans: we must view humans as other parts of the world’s causal web, that can be modelled analogously to how we model measuring instruments. This move is what allows the dispositionalist manoeuvre to be applied to theoretical perspectives, in addition to instruments. But then this would seem to undermine the idea that perspectivists should treat facts about how target systems manifest within perspectives any differently to any other facts. That is, if we find a perspectival analysis of worldly systems plausible, and we are also committed to viewing humans naturalistically, as no different in principle from other subjects of scientific investigation, then obviously we will favour a perspectivist analysis of humans and their cognitive constructions. This is not to say that the perspectivist is correct, but it is to say that the dispositionalist has given us no reason to reject perspectivism.

Ironically, a good illustration of this point is provided by the Votsis quote given above. “From the perspective of a normal human trichromat, the sky appears blue.” This is, Votsis says, a non-perspectival truth; it remains true from everyone else’s perspective too. Let’s consider this case in a bit more detail:

– It's not clear that a dichromat could come to understand the concept *blue*, just as trichromats cannot comprehend what tetrachromatic vision would be like. Consider: "From the perspective of a normal owl tetrachromat, the sky appears X." How can we understand this claim? Clearly not in terms of "X-ness", of which we have no experience. We would probably proceed as follows: take various items that all appear sky blue to humans. Then we can say that, in the owl's tetrachromatic visual system, items $i_1, i_2, i_3 \dots$ appear the same colour as the sky, while $i^*_1, i^*_2, i^*_3 \dots$ appear a different colour to the sky. This gives us some idea of how the owl sees the world. There are certain regularities in the world to which we are not sensitive, but which owls are.

– There are plenty of normal human trichromats who do not represent the sky as blue, since there are many societies that do not have any concept corresponding to our concept of *blue*. Many languages take what we call blue and green to be shades of the same colour; in other languages, what we call dark blue is grouped with black, and what we call light blue grouped with other colours (Kay and Maffi 1999). Votsis might try to get around this problem by saying that "appears blue" is simply describing a perceptual response that is shared by people across these cultures. A normal human trichromat in Ancient Greece would not call the sky "blue", and would not apply any other term corresponding to our concept *blue*, but she would have the same quale. Or if we want to avoid speculating about the experiential content of other minds: she would have the same physiological reactions in terms of stimulation of the retina, transmission of information down the optic nerve, and so on. But now we are no longer dealing with a proposition at all. "The sky appears blue" is a proposition, which is true from the perspective of a normal human trichromat. A person's experience of blueness, or the way the retina responds to particular wavelengths of light, are not propositions that can be true or false.

– What counts as a normal human trichromat? Among humans, there is a great deal of variation in colour perception. Hardin (1988: 76-82) provides a detailed examination of this. In medical and scientific contexts, normalcy in colour vision is defined according to various screening tests, whose results are not always in agreement. Furthermore, there are degrees of deficiency in colour discrimination, so the line between dichromacy and trichromacy, and the line between normal

trichromacy and deficient trichromacy, is not sharp. Another common variation is that the lens of the eye yellows with age, and this may happen at a higher rate among people nearer equatorial regions (Hardin 1988: 167). Since normalcy must be defined relative to a reference class, a person with yellowing that is statistically normal in one population may, per the same standard for normalcy, count as abnormal in another. In all these cases, it is up to us to draw the lines wherever we find the most convenient.

– The sky does not appear blue. It appears a variety of colours: black, pink, red, yellow, and grey are all common ones. Moreover, two normal observers directly viewing the same area of sky will perceive it differently if the surrounding environment is changed. Suppose one person is standing outside, and another person is viewing the sky through a hole in a wall, on which various other colours are painted. Paint the wall the right way, and the sky will appear almost any colour you like. The same colour will change its appearance as the colours surrounding it are varied. This is one source of many striking colour illusions. When we say that the sky is blue, we are engaging in idealization, simplifying away various other properties of the sky. What makes it true for us that the sky is blue are the special standards of evaluation that are brought with certain perspectives. As Teller might put it, it is true enough that the sky is blue, and the perspective supplies the standards for what counts as enough.

– What exactly is the sky? Obviously, the term does not pick out any determinate object. It might refer to anything above the surface of the Earth, or to a very loosely defined region of the Earth's atmosphere, or to a fictional "celestial sphere" surrounding the Earth onto which various colours and images are projected.

"From the perspective of a normal human trichromat, the sky appears blue" – this is true. But I don't think the perspectivist has any trouble at all in insisting that it is true only from particular perspectives. It is a perspective that uses our colour lexicon, that is informed by colour science, that categorizes human visual systems in a particular way, and that distinguishes a certain area above the Earth. The perspectivist should grant that there are dispositional facts, facts about how certain properties manifest under certain conditions, including within certain perspectives. These are also

perspectival facts. As for the categorical facts, such as facts about surface reflectances, we can take the same view. The sky appears differently under different perspectives, but we can describe the property of the sky that disposes it to appear differently. But what is this property? The perspectivist will say that this will also manifest differently under different perspectives.

9.2.3. *Behind-the-schemes*

A third approach to Escape from Perspective is Button's "behind-the-schemes" argument (Button 2013: 201-206). Button provides several illustrations of this, one of which we have already encountered: Putnam's world. We saw that when describing the alternative object-frameworks for Putnam's world, we took as fixed the claim that the world contains three atoms. The atoms stand "behind the object schemes", and when describing Putnam's world in terms of atoms rather than objects, we appear to have given a nonperspectival description of it. Certainly, we have stepped outside the different perspectives given by our object-frameworks. Button adds the point that if we assume that, say, compositional universalism and compositional nihilism are adequate frameworks for describing the parts of Putnam's world, then since universalism carves out the same parts as nihilism, plus others, Putnam's world must have more parts than is recognized by the nihilist. The nihilist can be treated as speaking with restricted quantifiers (2013: 203).

In my discussion of Putnam's world, I argued that it does not support an interesting kind of perspectivism. Then I suggested ways in which investigation into the real world is relevantly different. Button's objection is that regardless of the way the world is, and no matter how we specify alternative schemes, there must always be some way of carving out entities that is not relative to those schemes. Essentially, the issue is that we cannot describe alternative perspectives without also describing some shared world that they are perspectives on. If I claim that Newtonian mechanics and general relativity can be used to develop alternative perspectives on the solar system, then I must have some independent description of the solar system to identify it as the target that is treated by both perspectives. Without this, we could only say that there is a Newtonian solar system and an Einsteinian solar system, and there would be no particular relation between them. This seems absurd. But then if

there are true perspective-independent descriptions of things, this amounts to an Escape from Perspective. Anytime we try to articulate rival perspectives, we find ourselves with a perspective-independent description. As Button puts it, in terms of conceptual schemes:

From whatever perspective it is right to affirm that two rival schemes both succeed in talking about certain objects, that perspective countenances both kinds of objects without relativizing them to the two rival schemes. Otherwise put: the behind-the-schemes argument can be run *within* any conceptual scheme, and so conceptual relativism cannot be affirmed from within any conceptual scheme. But the conceptual relativist herself denies the existence of any scheme-transcendent perspective. (2013: 208)

First, a logical point. Caution is required in stating Button's objection, because under a natural interpretation, it would involve a straightforward quantifier-shift fallacy. Suppose we agree that for every instance of two rival perspectives on a given target, there is some further description of that target that is independent of those two perspectives. This does not entail that there is a single description of the target that is independent of all perspectives. It may be that perspective P and P* presuppose description D, while perspectives P1 and P1* presuppose a different description D1. Then a perspectivist will naturally take D and D1 to be rival perspectives, and in elaborating these perspectives she will appeal to some further description D2. So, if the behind-the-schemes objection is taken to have the conclusion that there must be a single, true, perspective-independent description of any given system, then it rests on a quantifier-shift fallacy. But I don't think that this is exactly Button's point. I take Button to be pointing to a potential incoherence in perspectivism: any attempt to state an example of perspectivism will have Escape from Perspective built into it, as it were, in virtue of the fact that our description must assume some perspective-independent facts. At the very least, perspectivism seems self-undermining in this way. We can imagine the dialogue between a perspectivist and her opponents going something like this:

Perspectivist: There are no perspective-independent facts.

Anti-Perspectivist: What does it mean for a fact to be perspective-dependent?

P: Consider Putnam's world, a world containing just three atoms... [at this point the perspectivist gives the example of different object-frameworks]

AP: Okay, but these different object-frameworks seem like alternative descriptions of the same fact: namely, that there are three atoms. Moreover, even if I accept that facts about the number of objects are perspective-dependent, what justification is there for supposing that there are *no* perspective-independent facts? To talk about perspectival facts, you have given a description that is independent of the perspectives: again, that there are three atoms.

The problem is that any example that is given of perspectival facts will always turn out to be grounded in a perspective-independent description. The perspectivist does not seem able to properly state her position. Now, it would be tempting to say that this problem goes both ways. After all, just as the perspectivist will always give some perspective-independent description, so that perspective-independent description is always vulnerable to the proposal of some alternative description. The number of objects in my room is a perspectival fact, depending on what object-framework I adopt. All the object-frameworks are applied to the same set of things: a chair, a bed, a cup, a computer, etc. So, "my room contains a chair, a bed, a cup..." is the perspective-independent description. But now the perspectivist will point out that there are plenty of other ways of carving up my room. The description is only perspective-independent with respect to the object/parthood perspectives. So we might think that both the perspectivist and the traditional realist are in the same boat. Any time the traditional realist tries to give a perspective-independent description, the perspectivist can give an alternative proposal.

I think that this is too quick. The realist's point here is that we have a general reason for thinking that there will always be the opportunity to Escape from Perspective, in any specific case. This is because any example of rival perspectives *must* provide a shared description of what they are perspectives on. This description will be provided by the perspectivist herself. By contrast, the perspectivist can hope that for any

shared description, there is a rival description that we would also take to be true in some circumstances, but no reason is given for thinking that there must *always* be such a rival. Moreover, were the perspectivist to ask the realist to state a perspective-independent fact, the realist would not necessarily commit herself to such rival descriptions. The realist can just say, for instance: “Kane’s room contains fifty objects.” She may be wrong that this is perspective-independent, but the example does not itself contain a rival perspective. So while the realist can say that there will always be a shared description, the best the perspectivist can do is point to past successful articulating rival perspectives. Many of these will be open to doubt: alternative categorizations of objects arguably amount to a mere verbal distinction, while the natural response to cases such as Newtonian vs Einsteinian perspectives is just to say that the latter is true, or at least closer to the truth. So some proposed rival perspectives turn out not to be genuine rivals, and in some cases we can straightforwardly favour one over the other, as a description of reality.

Bear in mind, however, that there *are* general arguments in favour of perspectival facts. We have already examined one of these arguments in much detail: the failure of fit argument. This argument aims to provide a quite general reason for favouring perspectivism. Whether perspectivists can in any specific case provide an example of a rival perspective is tangential to the success of this argument. Still, this leaves us with the challenge that there is a problem stating perspectivism, since any statement of it will involve appealing to shared descriptions. This is a genuine difference between the perspectivist and the traditional realist. I don’t see this as a particularly serious difficulty, though. It’s important to distinguish the statement of a position from an illustration of it. The point of Putnam’s world, and similar examples, is simply to illustrate what perspectival facts are. It may well be the case that to illustrate this, we must take certain objects as fixed. For this reason, a global perspectivism may not be “visualizable”, as it were, since in any particular case, there will be some aspect of our model that we do not at that moment represent as perspective-dependent. In the context of visualizing the conflict between rival perspectives P1 and P2, we take description D as fixed. This does not stop us from recognizing that there may also be rivals to D, even if we are not currently representing them. The situation here is similar to that of Moore’s paradox:

(M) It is raining, and I believe that it is not raining.

To assert (M) would be to imply a contradiction, since if I believe (M), then I am committed to the first conjunct, in which case I believe that it is raining, contradicting the second conjunct. I cannot coherently imagine a situation in which I would assert (M). Nevertheless, (M) might still be true: it is obviously possible that it is raining, while I believe incorrectly that it is not raining. Indeed, I have extremely good reason for thinking that there are plenty of “Moorean propositions” such as (M) that are true, since I know that all people have some false beliefs, and there is special nothing about me that would make me impervious to false belief. If somebody were to request an example of a true Moorean proposition, I could not give one – indeed, to suggest an example would involve me in an incoherence – yet I know that there must be some true Moorean propositions.

When I describe the rivalry between two perspectives, I must “step outside” those perspectives. I am no longer viewing the world from either P1 or P2; rather, I am viewing P1 and P2. Of course, this looks like “Escape from Perspective”. But if the perspectivist is right, this viewing of perspectives must itself be done from some other perspective. This other perspective will go unstated because we are *using* the perspective, not describing it; and in general, when using a perspective, you do not usually represent your own representations as perspectival. For example, if we are talking about the sky in an everyday context, I will endorse the proposition:

(B) The sky is blue.

As a perspectivist I will say that this is true relative to a particular colour perspective. So I will also endorse:

(B*) “The sky is blue” is true relative to perspective P.

But (B) and (B*) are, of course, different propositions. Even if I endorse (B*), I need not suppose that (B) is merely an abbreviation of (B*), or anything along those lines. If somebody asks me the colour of the sky then, in most contexts, I will simply talk about the sky: this will be done from a perspective that we hopefully share, but I

need not explicitly represent that perspective. Perspectives provide the means by which we represent the world; they are not usually part of the content of our representations. This is perhaps analogous to how light is one part of the means by which we perceive objects, but it is not itself something that we perceive (at least under normal circumstances). Since I will always be using some perspective or other to represent the world, there will always be some part of my description of things that is simply stated, without qualification.

One way to make this point is to consider from what perspective we might explicitly endorse a claim such as (B*). I take it that (B*) could be defended from the perspective of contemporary colour science. In this case, we would fill in the details of “perspective P” by filling in the details of human visual processing. Notice however that although colour science classifies objects, and connects those classifications to our social colour categories, it does not classify objects under the description of any particular social colour categories. Instead, it classifies objects in terms of, for example, surface spectral reflectances. After all, different societies divide up colours in different ways; from the point of view of colour science, no particular social colour categories are privileged. The classification scheme of the Himba is different to ours – they do not recognize anything corresponding to our “blue”, for instance – and colour scientists will simply try to describe and explain these differences.

To summarize: to give a description of anything, I must use some perspective P. I cannot simultaneously use P and describe P, because using P simply involves seeing things in terms of P, while describing P requires me to “step outside” of P. To describe P, I must adopt some other perspective P1, but again, I cannot simultaneously use P1 and describe. This is why any example of rival perspectives will take certain things as fixed.

9.3. A general objection

I have examined several ways of making the Escape from Perspective objection and have found that all are unsuccessful against the form of perspectivism developed in this thesis. I will now turn to a general objection against any Escape from Perspective argument.

All approaches to Escape from Perspective work by attempting to identify some underlying non-perspectival facts. There are certain propositions that are taken to be true, non-perspectivally; and then the apparently perspectival facts are understood as grounded in the non-perspectival facts. For example, while “the sky is blue” is true relative to a certain visual perspective, there are non-perspectival facts about the properties of the sky, and non-perspectival facts about the properties of the visual perspective, that ground the perspectival fact. The problem with this kind of move is that it undermines the justification for accepting realism in the first place.

We should begin by bearing in mind an important point raised by Ruyant (2020: 20). Ruyant notes that when perspectivists argue that some representation is perspectival, realists often respond by shifting the target of the representation. We have an initial representation of some target system *X*. Reasons are given for treating this representation as perspectival. Then the realist will note that there is a non-perspectival fact concerning *X* plus the users of the representation of *X*. “The sky is blue” is a perspectival representation, but that the sky is represented as blue within a particular colour perspective is not. But even if this is correct, it remains the case that the initial representation, our representation of the sky, is perspectival. Although Ruyant is concerned with representations, the same point can be made concerning perspectival facts. The traditional realist must shift from talking about some target system *X*, to talking about how *X* appears from a particular point of view, or how *X* manifests itself in particular conditions. The Escape from Perspective is achieved by moving up a level, as it were. There is no Escape from Perspective for *X* in itself. There are non-perspectival facts only concerning how *X* manifests in particular perspectives.

What role do such non-perspectival facts play in science? It seems to me that such non-perspectival facts would rarely be of much use. Part of the reason for this is that to represent such facts, we need to be able to specify the details of a given perspective: for example, the Escape from Perspective in the case of “the sky is blue” involves providing a model of the human visual system. But in general, it is extremely difficult to provide such details. Certainly, we rarely, if ever, find representations of perspectives in the sciences. The human visual system is only

part of a visual perspective, which involves a host of other elements, such as a particular cognitive framework for organizing visual perceptions. Recall Massimi's definition of a scientific perspective as "the actual – historically and intellectually situated – scientific practice of any real scientific community at any given historical time." Or here is Giere on maps:

I would like to say that the cultural background, the conventions for mapmaking, the designation of the region mapped, the specification of what features are mapped, and the degree of accuracy all determine a *perspective* from which the region is mapped. Every map reflects a perspective on the region mapped, a perspective built in by the mapmakers. In short, mapmaking and map using are *perspectival*. (2006: 75)

When we make or use maps, sometimes we do explicitly describe aspects of the perspective from which the map is constructed. For a 2D map of the globe, cartographers often have reason to draw attention to the method of projecting from a 3D Earth onto a 2D plane. The mathematical function that transforms geographical coordinates into Cartesian coordinates is one part of the perspective, and this is stated explicitly. But as is clear from Giere's quote, mapping involves a great deal more than this. For one thing, we are not really projecting the Earth onto a plane, but an idealized model of the Earth. The Earth could be treated as a sphere, an oblate spheroid, a geoid, or many other shapes. Particular features must be selected, the surface must be divided up in particular ways, some standard for accuracy and detail must be decided, and so on. For those using the map, much of this will be left implicit. This is of course the point made in the previous section that when using a perspective to represent a target, your representation does not itself represent the perspective used to generate the representation. We can step outside the perspective, and represent the perspective or parts of the perspective, as when we describe the mathematical function used in a map projection. But in many cases, it is difficult if not impossible for us to describe perspectives in any detail. In some cases, the relevant information may simply be unavailable: the evidence concerning the motivation and background assumptions of a particular scientist or research group may never have been recorded.

I claim that facts about perspectives are often inaccessible to us. Escape from Perspective involves showing that perspectival facts are grounded in non-perspectival facts. But these non-perspectival facts include facts about perspectives, and so often, these non-perspectival facts cannot be stated. In itself, this is not necessarily a problem. The mere fact that we cannot state such facts does not entail that there are no facts. However, it raises trouble for the traditional realist, given how their position is usually justified. Why should we believe that our best theories are true? As we have seen in previous chapters, the usual argument appears to the remarkable success of those theories. Recall that realists generally endorse a form of inference that Kitcher (2002: 348) calls the “success-to-truth” rule:

S plays a crucial role in a systematic practice of fine-grained prediction and intervention.

So, S is approximately true.

This much can be accepted by a perspectivist too, since perspectivists will also hold that various scientific claims are true. It’s just that they are true only relative to some perspective. Now, one of the standard challenges to realism is that there appear to be a host of counterexamples to the success-to-truth rule. Over the history of science, we find a variety of theories that were successful, but that are no longer taken to be true (Laudan 1981), and even our best contemporary theories are littered with idealizations. This challenge has led to selective realism: most realists do not claim simply that our best theories are true, but that we can identify specific parts of theories that we are justified in believing. Indeed, this selective skepticism is found in Kitcher’s statements of the success-to-truth rule: S must play a *crucial* role in a *systematic* practice... The luminiferous ether was postulated as part of a remarkably powerful theory, but it did not play a crucial role in the derivation of the predictions of that theory – or so realists will say (e.g. Psillos 1999). Most realists defend their position by trying to separate the wheat from the chaff.

Notice that the same kind of move is being made in the application of Escape from Perspective. For the traditional realist, perspectival truths are in some sense inferior, so that we are justified in taking them as *truths* only because they are grounded in non-perspectival truths; or perhaps the very notion of “perspectival truth” is just

incoherent, so there are no perspectival truths in any sense – whatever the case may be, perspectival truths are the chaff, and what we are really after are the non-perspectival truths. A reductionist might similarly say of “higher level” propositions that either they are grounded in the appropriate “lower level” propositions, or there simply are no higher-level truths: e.g. propositional attitudes reduce to neural states, or propositional attitudes do not exist. The traditional realist is a kind of reductionist with respect to the perspectival propositions. Now, the problem with this is that, as I have suggested above, the relevant non-perspectival truths in question are often not part of the content of any scientific theories. Take:

(B) “The sky is blue” is true relative to the normal trichromatic human visual perspective.

So “the sky is blue”, if it is true, is a perspectival truth. This truth can be grounded in the non-perspectival truth of (B) plus non-perspectival truths concerning the human visual perspective. In this case, we do have successful scientific theories describing aspects of this perspective. Contemporary colour science has provided very useful models of parts of the human visual system, such as opponent-process theory. That is, various claims describing the human visual perspective do play “a crucial role in a systematic practice of fine-grained prediction and intervention.” But now consider:

(O) Opponent-process theory is true relative to the perspective of contemporary colour science.

What exactly is the perspective of contemporary colour science? We can say something about this, and this can also be the subject of scientific study. Sociologists and anthropologists may develop models of contemporary colour science. But it’s quite clear that claims about this scientific perspective do not play a crucial role in theories that support fine-grained prediction and intervention. Moreover, there are reasons to think that such theories are necessarily inaccessible. Suppose we had a theory that allowed us to predict the future course of contemporary colour science. At best, this would have to be limited purely to phenomena such as the institutional structure of the science. We could not predict the content of future colour science perspectives. The reason for this is that, as Popper (1960) has pointed out, any

reliable prediction of the content of future science would amount to a discovery of that content. That is, the prediction that future scientists will come to know some theory T would itself entail that we already know T. If I know today that tomorrow I will know that my brother stole my chocolate cake, I must *already* know that my brother stole my chocolate cake.

Clearly, not everything about the growth of knowledge is beyond the possibility of reliable prediction. Perhaps we can predict that the existence or non-existence of Planet Nine will be conclusively confirmed within the next ten years. Or we might predict, on the basis of a pessimistic induction, that current theories are likely to be displaced. But we cannot predict what anomalies will drive the theoretical changes, nor what specific theories will be adopted in place of current ones. In colour science, opponent-process theory is the best we can do; and if we somehow knew what problems it would face in the future, and we knew that theory T would solve those problems, then we would already accept T instead. In any case, even if Popper's point is too strong, it's quite clear that in practice, we do not have anything even approaching a sophisticated, empirically successful model of most scientific perspectives.

So the problem for the traditional realist is this. Realists endorse the success-to-truth rule: we are justified in believing that propositions are true when those propositions play the right kind of role in theories that exhibit true novel predictions, broad explanatory scope, applications in controlling phenomena, etc. Now when the realist appeals to Escape from Perspective, she will try to show how the proposed perspectival truths are grounded in non-perspectival truths. These non-perspectival truths must include truths about scientific perspectives. But we do not yet have, and in some respects perhaps could never have, successful scientific models of scientific perspectives. Descriptions of scientific perspectives do not "play a crucial role in a systematic practice of fine-grained prediction and intervention." In which case, for the realist, we are not justified in taking such descriptions to be true at all, let alone non-perspectivally true. Of course, we can always reduce the force of this problem by holding that the success-to-truth rule is merely sufficient for justifying belief, and that there are other forms of justification available, so that claims about scientific perspectives can be justified. But then the easier we make it to justify theoretical

claims, the greater will be the number of the apparent counterexamples to realism: theoretical claims that scientists were justified in believing in the past, but that turned out to be false. While perspectivism is sometimes presented as a limited form of realism, my discussion suggests a different approach: it is the only way of preserving realism with any justification.

There is a perhaps more general point to be made here. For philosophers attracted to scientific realism, it seems natural to say that the best explanations, and the theories that we can take to be justified, are those produced by our best science. The realist intuition is that science is our most powerful tool for getting at the way the world is, at least in certain domains, and so in those domains the epistemically responsible person will let her beliefs be guided by the best science. Yet in practice, this is not how philosophically sophisticated realists tend to proceed. It is not science that guides their belief, but the philosopher's "cleaned-up" science. We find this in the selective realist's attempt to quarantine particular aspects of theories in response to problems such as the pessimistic induction and idealization; we see this in Quine's "regimentation" of theory, where theories must be translated into logically rigorous language and ontological commitments minimized (Quine 1960: 157). We also see this in *Escape from Perspectivism* response to the perspectivist: we are justified in believing those propositions that have been appropriately grounded in non-perspectival facts. But it is not scientists who show that this grounding obtains. In all of these cases, we must go beyond what is given in scientific theories, and we must be applying norms of theory-construction that are somewhat at odds with those that govern scientific practice. If the success of science is what it is impressive, it is unclear why the philosopher's regimented science would inherit this success.

9.4. Incoherence?

The *Escape from Perspective* argument has been pressed most forcefully by Chakravartty. Ultimately, I do not think it is successful. Recently however, Chakravartty has presented a different anti-perspectivist argument, and one which connects with concerns about lexicons and failure of perfect fit. I think this argument can be dealt with more quickly, but it is worth detailing. The argument again takes the form of a dilemma. Chakravartty objects that the perspectivist argument

appealing to different linguistic frameworks with which to describe the world – this is essentially the failure of fit argument – faces a similar problem as the appeal to inconsistent models: it will turn out that these frameworks are merely instrumental, or that our worldview is simply incoherent (2017b: 185). Essentially, the difficulty is this. On the one hand, we might take it that our different ways of describing the world have no ontological significance. In that case, we have adopted an ontologically deflationary stance, where these descriptions are treated instrumentally. On the other hand, if we suppose that these descriptions are true, and that the entities and relations described really exist, we will be committed to an incoherent, contradictory world.

Chakravartty targets the perspectivist claim that there are no perspective-transcendent facts. He takes it that there are two ways of interpreting this claim. On the first horn of this dilemma, “Talk of ontology here is simply a way of talking about one’s – or a scientific community’s – choice of paradigm or framework for describing the world” (2017b: 185). We cannot know about the nature of the world itself, but only about our frameworks for describing or representing the world. There is no room for perspectival facts here, nor even just perspectival knowledge. There are various frameworks we might use for organizing experience, but these have no ontological significance so do not give us knowledge of anything beyond experience. Chakravartty compares this to the Kantian position that we know only the phenomena, while the noumenal world is forever hidden: “our knowledge of the world is limited to the phenomenal world, the world of our experience; it is not knowledge of the world in itself” (2017b: 185).

I think that Chakravartty moves too quickly here. There are a couple of ways of interpreting the position he outlines. We might take it, first, as asserting knowledge of some circumscribed class of things – say, of observables or perhaps just experience. Then the position is instrumentalist about claims beyond this. So we take our descriptions of the world as being useful tools for systematizing experience, say, which might then support prediction and control. Yet this position requires there to be “uninterpreted experience” or “uninterpreted observables”, to which different frameworks are applied. Naturally, such a position is unattractive to many philosophers. Few of those inclined towards perspectivism would accept anything

like this. A perspectivist will want to give a perspectival account of knowledge in general – knowledge of observables and, if there be any, knowledge of unobservables.

Alternatively, taking a more Kantian line, we might hold that we have knowledge beyond our immediate experience, or beyond observables, but that this knowledge is of something more like Kant's phenomenal world. That is, the fact that we can know the world beyond what is immediately experienced does not entail that we can know the world "in itself", or even that there is such a thing as what the world is like "in itself". This seems to be more in line with perspectivism. Of course, the perspectivist will allow for different perspectives and changing perspectives over time. But why should we not think of claims generated within such perspectives as describing facts? For example, suppose I make a claim about the length of the coastline. This is dependent on my choice of measuring units. This is, in a straightforward way, a claim about the world, and there is no reason to treat this as merely instrumental, or as "deflationary" about ontology. I am straightforwardly attributing a particular property (a particular length) to a thing in the world (namely, a coastline). Indeed, we can't reframe this as merely a useful instrument for talking about observables, because the coastline is already an observable. If we adopt an instrumentalist interpretation of my coastline-judgement, presumably we would have to assert only that it is a useful instrument for systematizing and predicting sensory experience, or something along those lines. But now, again, it seems we are appealing to some notion of uninterpreted experience.

Now, however, we face the other horn of the dilemma: incoherence. Chakravartty says: "For the neo-Kantian perspectivist ... there is no fixed world to know. ... In virtue of conceptualizing the world in conflicting ways, scientists would somehow create a fundamentally conflicted reality" (2017: 186) To attribute incompatible states of affairs to the world, and where facts in the world change depending on one's perspective, is, Chakravartty says, simply incoherent. We would be supposing that the world itself somehow bears contradictory properties. In earlier work, Chakravartty calls this a "Frankenstein" world (2010: 411). Appealing to different scientific practices, different investigative goals, different classification schemes, and other components of scientific perspectives, does not help. We simply cannot take fluids,

for instance, to be “both continuous media and collections of discrete particles but in different contexts, because in both the sciences and in everyday life, we take a fluid to be the same thing across different contexts of investigation” (2017b: 186).

In supposing that the same system can have incompatible properties, the perspectivist is committed to a prima facie incoherence: Liquids are both continuous media and collections of particles. Perspectivists try to avoid this through relativization: from perspective P, system S has property A; from perspective P*, system S does not have property A. But the relativization fails, according to Chakravartty, because we may assume that the system is the same thing no matter what perspective we are taking on it. Systems and objects are fixed across different investigations. The water that comes from my tap is the same thing, whether I’m using it to quench my thirst, or whether a chemist is using it to create a caustic soda solution for the purpose of an experiment. Indeed, many perspectivists do not wish to deny this: Giere, at least, accepts as a methodological assumption that the world is unique. If anything, this makes the incoherence charge only more pressing. Not only must we make sense of the notion that facts are different from different perspectives, but apparently we must also have it that the facts are the same across perspectives. There must be some sense in which there is identity across perspectives, because we want to say that different perspectives can be perspectives on the same system.

I have already suggested that Chakravartty does not acknowledge all the available positions. His dilemma is a false dilemma. To make this clear, Chakravartty argues that the perspectivist turns out to be committed to either:

(I) Instrumentalism: We have knowledge of sense experience or observable objects, but nothing beyond that. Alternative models and alternatives means of classifications are merely useful tools.

Or

(II) Profligacy: The world itself has contradictory features. Water is literally both a continuous fluid and a collection of discrete particles.

In contrast to both, I have already suggested:

(III) We have knowledge of the world beyond experience/observable, but this does not constitute knowledge of the world in itself; we have knowledge of perspective-dependent facts.

Facts are true propositions – recall that Chakravartty himself accepts this analysis of facts. The argument is then that whether a proposition is true is dependent on perspectives. So facts are perspective-dependent. Consider again the coastline example: in reporting alternative measurements, is there actually an inconsistency here? Only if we assume the use of the same measurement system. I say that the coastline is X miles long. This is a claim about the world, but a claim about the world from a particular perspective: its truth is dependent on that perspective. The length of the coastline is a perspectival fact. Of course, if we insist on “detaching” facts from perspectives, then if we maintain that there are various alternative lengths of the same object, we apparently face instrumentalism or profligacy. But I see no reason why we must think of facts as “detached” in this way. At least, this is something that those who object to perspectivism must argue for. But to give such an argument would itself be to deny perspectivism. That is, anybody arguing that there are perspective-independent facts is already arguing for the denial of perspectivism. In that case, Chakravartty’s dilemma becomes dialectically redundant. The dilemma can avoid begging the question only if it is supported by some further argument that would itself, if successful, refute perspectivism anyway.

Chapter 10

Toward a New Perspectivism

10.1. Introduction

In Chapter 1 of this thesis, I noted the variety of different approaches to perspectivism that have been explored in the literature. This thesis contributes to this diversity. I have used Giere's work as a springboard but have developed a view that no doubt differs in many ways from his own: in particular, my perspectivism lies more on the antirealist, constructivist side of the spectrum. Given its emphasis on classification and kinds, we might call it "theoretical kinds perspectivism" (TKP). Shortly, I will provide a summary of TKP, relating it to the six features of perspectivism that I distinguished in Chapter 1. Before doing this, however, I will note a potential worry about this project. We might wonder about the philosophical value of putting another form of perspectivism on the table, when there is much work to be done in clarifying and applying the perspectivisms currently available. I have three points in response to this concern.

First, some of the different forms of perspectivism are addressing different questions, even if the underlying motivations of the views are similar. I have examined a metaphysical form of perspectivism – "metaphysical" in the sense that it is about what it is that we have knowledge of, not the epistemological question of how we form knowledge, which is the focus of, for example, Massimi's perspectivism. Bearing in mind this diversity of goals, it is to be expected that various positions may be developed; and they may not, in the end, all be incompatible. I rejected Massimi's approach to perspectival truth, but this was in the context of working out a distinctively perspectivist position with respect to a question that Massimi was not concerned with. An analogue to Massimi's epistemological perspectivism may well be consistent with TKP. This may be an interesting question for later research.

The second point is that, if my arguments are on the right track, then I have made some progress in trimming the thicket of available views. Chapters 2 and 3 argued that the prospects for instrumental perspectivism are poor, at least if it is treated as a separate position. It is defensible, but only as a consequence of a theoretical

perspectivist account of instrumental models. Chapter 4 argued against the attempts to develop perspectivism by replacing the concept of truth with alternative representational concepts such as similarity; and also argued that perspectivists should embrace relativism. This is all within the framework of developing a metaphysical form of perspectivism. Regardless of whether such a view is held to be plausible, the arguments of this thesis do clear out some of the conceptual jungle, and identify a position which is coherent and distinct from the traditional realist and antirealist options.

Finally, although the details of TKP differ from Giere's position, I think that it does constitute a *perspectivist* position, as Giere understood that term. I have argued that we have access only to perspectival facts, where facts are understood in terms of true propositions. Already, of course, this framing differs from Giere, who instead emphasizes models and similarity; again, Chapter 4 outlines the reasons for resisting this approach. But as we have also seen, Giere often makes claims that are naturally interpreted as affirming perspectival facts. This, of course, leads to the question: What exactly is a perspectival fact? And this requires an understanding of "perspective". In Chapter 1, I outlined how drawing from Giere's colour analogy, we can specify six features of perspectives. I shall now return to these six features, relating them to the position developed in this thesis, to demonstrate how TKP constitutes a form of perspectivism in Giere's sense. This also serves as a useful summary of the position. These six features are:

- (a) Facts are a product of the interaction between perspectives and the world.
- (b) Claims made within a perspective may be true or false.
- (c) Perspectives track regularities in the world.
- (d) Perspectives are partial.
- (e) Different perspectives are not necessarily incorrect/inaccurate.
- (f) Different perspectives are better or worse depending on one's purposes.

I will now elaborate on these points.

10.2. *Interactionism*

Colours do not exist independently of visual perspectives. Colour properties are a product of interaction between organism and environment. Our visual perspectives are determined partly by our physiology: the response of the retina to particular patterns of light, the transmission of information down the optic nerve, and so on. But these perspectives are also constituted by our conceptual understanding. Not all normal human trichromats will carve up colour space in the same way; the distinctions between colours are partly socially constructed. Colour perspectives are not just visual, but also conceptual. Scientific perspectivism attempts to generalize this interactionist idea, drawing on the conceptual contribution of perspectives. This is expressed in the perspectivist's commitment to the view that science gives us knowledge only of perspectival facts, where to quote Chakravartty (2010), "a perspectival fact is a proposition that is only true from, or within, or relative to, a given perspective (or limited set thereof)."

How does TKP capture the interactionist idea? To summarize, perspectives are partly constituted by the means of classification, and there is no perspective-independent standpoint from which such perspectives or their schemes can be evaluated. Objects are constituents of the world; properties are features of objects, abstracted from the objects that bear them. Any description of objects, properties, their relations, and so on, must be made from within a perspective, using its classification scheme. Crucially, our theoretical lexicons cannot be taken to mirror any lexicon of the world itself, or natural kinds. The divisions in nature are drawn by us. This is the central perspectivist argument, developed across chapters 5 to 8. Just as colour properties are a product of interaction between organism – with its perceptual system and its conceptual resources for interpreting perceptual input – and environment; the same is true for all other properties.

10.3. *Truth-value*

Claims about colour may still be true or false. For a normal human trichromat, living in English-speaking societies, it is true that grass is green. Speaking from within a perspective, and with a classification scheme in place, claims about the world can be

true or false. This expresses a realist element of TKP: there is an important sense, even within a perspectivist framework, in which we are getting things right or wrong. The fact that truth is evaluated relative to a perspective does not entail that there is no truth, or that “anything goes”. It is worth outlining the standard theories of truth, to demonstrate that there is nothing in them that is in tension with TKP. Indeed, some theories of truth are, if anything, more straightforwardly compatible with TKP than with traditional realist approaches (see Kirkham 1992 for a summary of theories of truth). Of course, this list is not exhaustive, and it is no doubt possible to propose some theory of truth on which perspectival truth is ruled out. But this is not the case for the standard theories.

10.3.1. *Correspondence theories*

Truth consists in the correspondence of a proposition to some state of affairs in the world. The arguments of this thesis have been implicitly framed in terms of correspondence theory, partly because it is still the most intuitive account of truth, and partly because, of all the traditional theories, this is theory that, I think, at least *prima facie* seems most problematic for the perspectivist. The reason is that whether a proposition corresponds to the world is, it seems, determined by *the way the world is*, not by our perspectives. If we claim that the truth consists in correspondence to the world, and that truth is in some sense constructed within perspectives, then it seems we are committed to an ontological plurality of worlds along the lines of Goodman (1978). This view has its own problems, but what I want to note here is that, while Goodman may have similar motivations to perspectivism, for a perspectivist to embrace Goodman’s worldmaking would seem to amount to giving up perspectivism. The perspectivist claims that there can be multiple legitimate perspectives *on the same target*; Goodman simply denies that there can be a single target for adequate but incompatible theories.

More precisely, Goodman holds that there are true conflicting “versions” – where a version is something similar to a perspective – and that to each true versions there answers a different world. Versions are not versions of the same target system. The reason for introducing different worlds in the first place is that if different, incompatible versions are true of different worlds, there is no longer any conflict

between them. V1 and V2 are incompatible. but V1 is true of W1, while V2 is true of W2. So we have straightforward correspondence between V1 and W1, and V2 and W2. V1 and V2 are straightforwardly, non-perspectivally, true, but each is true of different worlds.

There is, however, an easy fix that allows us to combine perspectivism and correspondence theory, and there is good independent motivation for this fix. Goodman says there are true conflicting versions, and to each true version answers a world. But what if there are no true versions? As perspectivists often like to emphasize, even our best theories involve a host of idealizations and simplifications, and all of them are incomplete in various ways. Realists answer this by introducing approximate truth. If we talk of approximate truth, the correspondence argument for many worlds no longer goes through. Approximate truth is just a special kind of falsehood. So it is not clear how approximately true versions could make worlds. After all, if a version has a world corresponding to it, then that version, surely, is just true. That is what it is to be true, per the correspondence theory: to have a corresponding world.

What is it for a match to obtain between a proposition and the world? The first thing to note here is that any description of the world will be incomplete and simplified. If we acknowledge that true statements can be made such as “the table is flat”, this is truth relative to a perspective that specifies certain standards of evaluation. That is, we must specify the relevant standards for what counts as a “match”: in the context of a precision experiment, the claim is false; in the context of a desk on which to place everyday objects, it’s true, since we have different requirements for what counts as flat (cf. Giere 2006: 87-88; Teller 2019). Then we must specify the components of the world in the comparison: as has been argued, there is no independent “lexicon of the world”. The point is that when we construct a model, we take it to correspond to the world; we have a model → world comparison. But the “world” here itself is another model; it is really a model → model comparison. The description of the world itself requires some perspective, some way of carving things up. Within perspectives, correspondence obtains; this is analogous to the point made in Section 8.6. that within perspectives, we may have essence kinds and homeostatic property cluster kinds.

10.3.2. Deflationist theories

The basic motivation of deflationism is the thought that truth is not a substantive property. To assert that a proposition is true is equivalent – in some important sense – to simply asserting the proposition. There are different forms of deflationism, but all take it that truth is defined in terms of the equivalence schema: “P” is true if and only if P. Beyond claims such as this, there is nothing further to be said about the nature of truth. No further content is added to a proposition by the assertion that the proposition is true.

This is straightforwardly compatible with TKP. Whether a proposition is assertable is dependent on one’s perspective. So whether it is true is dependent on one’s perspective. TKP need not suppose that there is more to evaluating ““P” is true” than is involved in evaluating “P”; just that in both cases, we necessarily appeal to the resources of some perspective or other, and there is no neutral standpoint from which such perspectives can be evaluated. This is all TKP needs. Of course, one worry here is that TKP takes it that truth is relative to perspective, and it may be thought that this involves a commitment to a more substantive theory of truth than is permitted by deflationism. However, we saw in Chapter 4 that the relativization need not be taken as part of the *content* of the claim ““P” is true” – that is, although the truth of P is relative to some perspective or other, the claim ““P” is true” need not be taken as abbreviating the proposition ““P” is true relative to perspective X.” So truth-claims can be taken in the straightforward, commonsensical way that the deflationist urges.

If there is a problem with deflationary approach, it is that they are too *permissive* about which propositions count as true. For some deflationists, the proliferation of truth is a feature, not a bug. But there is a genuine difficulty, from a deflationist point of view, in accommodating the notion of useful falsehoods, which obtain from almost every perspective. When a scientist models a star as if it were constituted by an ideal gas, it is recognized by all that this is simply false, yet arguably it is appropriate to assert this falsehood. Along the same lines, the deflationist runs into problems capturing *approximate* truth. While this is a problem for deflationism, it is not a

problem that results from TKP, since everybody recognizes such useful falsehoods. The key point for my purposes is that deflationists have no trouble counting various claims within perspectives as true.

10.3.3. Coherence theories

According to coherence theories, a proposition is true just in case it belongs to a maximally coherent set of propositions. There is much debate about what exactly constitutes a coherent set of propositions: consistency is a necessary condition; but in addition, we may insist that the propositions in question “fit well” together in some broader sense. So we may further virtues of standard virtues of simplicity, explanatory power, theoretical unification, pragmatic utility, etc., in assessing coherence. However we spell out what counts as coherence, the key point is that the truth-conditions of a proposition are not objective states of affairs in the world (at least not directly), but rather other propositions.

Coherence theories initially may seem to fit naturally with TKP. Indeed, a classic objection to coherentism as a theory of truth is that it entails relativism, since it seems that multiple inconsistent sets of propositions could nevertheless bear the relevant coherence relations. However, it could be argued that there is a snag in holding both coherentism and TKP. A crucial point of TKP is that alternative perspectives are easily available – indeed, we often shift from one perspective to another, depending on what inquiry we are engaging in. Perspectives can be very fine-grained, as we saw in the discussion on the construction of theoretical kinds: we may take different perspectives on the gulls around the Arctic, for instance. Now the worry about this is that it is incompatible with coherence: there is not coherence across different perspectives, so if, as TKP claims, I may hold multiple perspectives, then my beliefs will not be coherent overall.

Recall, however, that TKP claims that truth is constructed within perspectives; there is no non-perspectival truth, and no neutral standpoint from which perspectives are evaluated. Coherence can still obtain within perspectives, and this is all that TKP needs, from the point of view of coherence theory, to have perspectival truth. It is also worth noting that if we think of the production of knowledge as a social

enterprise, which involves offloading cognitive work onto other people, then traditional coherence theory will also face the same challenge. Even from traditional coherence theory, alternative coherent sets of propositions are potentially available, hence the relativism challenge. The traditional coherentist takes it that on an individual level, we can achieve maximal coherence. But there is no reason to expect a group of individuals all to arrive at the same coherent set of beliefs. So if knowledge is to be understood as the product of the group rather than the individual, we have the same problem of apparent failure of coherence due to incompatible perspectives. Again, however, it is enough to say that coherence, and hence truth, obtains within perspectives.

10.3.4. Pragmatist theories

The underlying intuition of these theories is that the true propositions are those which it is, in some sense, useful to believe. Pragmatists draw a tighter connection between truth and our epistemic practices; they begin with the question, what role is the concept “truth” in some given area of discourse? In general, the truths will be those propositions that play important roles in our best theories and that have fared well against critical examination. Different pragmatists will fill out these details in different ways. Some classic formulations of pragmatism are not hospitable to TKP. For Pierce, the truth is that which will be accepted at the “end of inquiry”, “the opinion which is fated to be ultimately agreed to by all who investigate”. This is incompatible with TKP’s pluralist and relativist commitments. However, Pierce makes a very strong claim here – stronger even than many traditional realists would accept. Most pragmatists do not take themselves to be committed to the notion of an end of inquiry, or to convergence on one single truth. The general pragmatist idea that truth has a close connection to utility – “the expedient in our way of thinking”, as James put it – is clearly in line with perspectivism. Indeed, if pragmatists have a problem, it is that, like deflationists, they may well be too permissive about what they count as true.

This discussion has, of course, been brief. But it should be clear that there is no prima facie problem with assessing truth and falsehood within a perspectivist

framework. Perspectival truth is compatible with a broad variety of truth theories. Within a perspective, we may still make claims that can be true or false.

10.4. Perspectives track regularities in the world.

Recall the example of red-green colour discrimination allowing organisms to pick out ripe fruits. Crucial to the interactionist idea is that perspectives do not float free; every useful perspective is a perspective on some aspect of the world. There are objective regularities in the world that our perspectives hook up to. This is a further element of the TKP's realist commitment. When a proposition is true within a perspective, then it plays some role in tracking features of the world. Of course, the danger here is Escape from Perspective. Once we describe the objective regularity, and we describe the perspective on that regularity and how these interact, it seems that we now have a non-perspectival description of the situation. I have addressed this challenge in the previous chapter, where I argued that three different ways of elaborating on this challenge fail, and that the attempt to identify non-perspectival facts is in tension with the realist's use of the success-to-truth inference.

But if it is, as it were, "perspectives all the way down", then in what sense can perspectives track objective, mind-independent facts? How can facts be perspective-dependent but mind-independent? After all, perspectives are partly constructed by us. While some aspects of perspectives might be thought of as "readymade" – think of, for instance, the structure of the visual system, or the relativity of motion to reference frames – all our perspectives are partly conceptual, and in the context of science, the product of scientific theorizing. This might seem to be a straightforward incoherence at the heart of perspectivism. However, this challenge is, I think, quite easily dealt with.

There are many senses of mind-independence. The relevant sense in this context, where we are claiming that perspectives track mind-independent features of the world, is *location*. That is, an object or event is mind-independent just in case it occurs outside of the mind. This distinction plays a crucial role in our best theories. This is clear if we start by thinking not in terms of our own view of the world, but in terms of how other people relate to the world. Kitcher (2001: 181) provides a useful

example to illustrate the point. He asks us to consider observing a subject who is using a map, trying to work out the best path to her destination. We take it that the subject has some sort of mental representation, or at least some sort of mental activity, prompted by viewing the map; and that this relates in some way to objects that are located outside of her mind. As Kitcher says:

Central to our ordinary explanation of what the subject does is the idea that she represents objects that would exist even if she were not present. We take the dots on the map as corresponding to things we can pick out in her environment (underground stations), and we think that the associated items in her mental state also correspond to those things. (2001: 181)

There is nothing controversial or even particularly informative about this description. Obviously, a proper explanation of the subject's behaviour would require much more detail. But already at this stage, we are drawing a distinction between minds and the external world, and this distinction is retained in the more sophisticated theories that fill out this sketch. In explaining the behaviour and their theorizing of others, we look at correspondence relations between their mental representations and the world. Perhaps, in the end, this is not a requirement. Perhaps there are powerful theories of such behaviour that do not require drawing such distinctions. We need not rule out this possibility, but it is quite clear that at present, our own theories are not of this nature. The best means we have for predicting and explaining the behaviour of others draws a mind/world distinction. Moreover, there are many contexts where the activities of a person's mind do not causally influence the events in the world. We can consider a subject who is not just reading a map, but who is performing various experiments, developing new concepts and theories, and so on, in order to predict and explain phenomena, say, the behaviour of a gas. Her conclusions need not make any difference to our predictions and explanations of the gas. Add other subjects, who are each constructing their own theories: still, we find that the behaviour of the gas remains invariant under these different conditions. So the gas is independent of the minds of the subjects.

To this, we need only add that there is nothing special about us. A requirement of psychological and sociological theories is that they apply to ourselves as well; since

we are all humans, since we relate to the world in similar ways, we can apply similar frameworks in predicting and explaining our own behaviour. When we construct perspectives on human interaction with the world, we don't treat our own position as privileged. So the notion of mind-independence in Kitcher's account, where objects in the world are independent of the minds of the people we are observing, is generalized, so that these objects are treated as independent of our minds also. The distinction between humans and mind-independent objects is given within our perspective. Thus, we correctly describe our theories and models as, very often, tracking mind-independent facts.

10.5. Perspectives are partial.

We saw in the case of colour perspectives that the visual system is responsive only to a particular range of electromagnetic radiation. In this respect, the visual system is partial. Along the same lines, no perspective can describe everything, and theories and models constructed within perspectives can capture only particular aspects of target system. As it turns out, partiality is a rather trivial property. Few philosophers would deny that even in our best theories have a restricted scope of application. However, the partiality of perspectives does bear an important connection to the interactionist feature. As noted by Giere (2006: 66), a model that achieves an exact fit would need to be a complete model. Those aspects of the system that the model does not represent, and its relation to other unrepresented systems, will have some influence on the aspects that the model does represent. Thus, theoretical kinds are a practical necessity. Moreover, the discussion of theoretical kinds shows how the partial focus of perspectives allows for the construction of theoretical kinds. Consider again the example of gulls around the Arctic: if we shift focus from the whole ring species to a single country where the ends of the ring meet, we create the conditions the application of the biological species criterion.

10.6. Different perspectives may both be accurate.

This is the relativist commitment of TKP. Claims about the world are always made from within a perspective; and we cannot step outside all perspectives. Different perspectives on the world may reveal different properties, constructed by the

interaction of those perspectives with the world. There is no independent “lexicon of the world” that could provide a neutral standpoint, or that perspectives could match or fail to match. Moreover, it is clear that in different contexts, different theoretical kinds will be appropriate.

It might be objected that TKP is committed to a global antirealism or skepticism. After all, all theories and models will involve idealization; and it has already been argued that the language of our theories does not track the language of the world. What I have called “theoretical kinds” are simply useful falsehoods, or something along those lines. Moreover, as we saw in the previous chapter, this can be generalized even to claims about perspectives: not just “the sky appears blue”, but even “from such-and-such perspective, the sky appears blue” is littered with idealizations and so, in some important sense, is not strictly speaking true. But this is just to say: it is not true.

This worry, however, simply overlooks what has already been said in defence of a notion of perspectival truth. It is worth noting, first, that realists are already committed only to the approximate truth of our best theories. That is, not even realists suppose that our theories are exactly, perfectly true; and as I suggested above, *approximate* truth is really a glorified kind of falsehood. The point is that all parties are working with a view of truth that does not require perfection. I have also argued, in section 10.3., that different theories of truth are, *prima facie*, entirely compatible with truth constituted within perspectives. Indeed, if anything, TKP has the upper hand here. Realists hold that our best theories are approximately true; but what exactly are the standards for what counts as approximate truth? How close to the truth must a claim be, to count as approximately true? It is hard to see how this could be fixed perspective-independently. These standards are our standards, and are dependent on the particular field of inquiry. TKP builds this into its account of truth explicitly.

10.7. Perspectives and purposes.

The relativist commitment of TKP is no obstacle to talking about the utility of different perspectives. We can recognize that some perspective is accurate within a particular context, while still judging the propositions stated from that perspective to be false

given the goals of our own inquiry. Of course, these judgments are, like any other judgment, made from within some perspective or other; the coherence of this relativist view has been discussed in Chapter 4. The important point is that perspectives are not totalizing; we need not see the world only from one perspective. In fact, this is probably rarely the case. The theoretical kinds that we construct are sensitive to specific details of our inquiry: how exactly we draw the lines of biological populations or biological individuals will change depending on whether we are studying an ecosystem over years or decades, for example. This distinguishes perspectives from worldviews or Kuhnian paradigms. Perspectives are more limited in their scope.

Perspectives must be evaluated not simply in terms of whether they correctly represent the world, but also based on whether they satisfy particular goals of the users of the perspective. Different perspectives are better or worse depending on one's purposes. This is a point that is widely accepted: there is more to knowledge production than simply getting at the truth. The standard view, as we saw in Chapters 7 and 8, is that our interests play a role insofar as they guide the *selection* of various possible perspectives. Since all our theories and models are partial, and none can represent any target system in all its detail, we must decide what properties we are interested in tracking. There is a plurality of kinds in the world, so a plurality of ways of "carving nature at its joints". For TKP, however, our theories deliver theoretical kinds, and such kinds are strictly speaking never instantiated. Interests play a role not just in the selection of which independently-correct classification scheme to adopt, but in the construction of the kinds that are classified within the perspective. The question of whether a given perspective provides accurate representations does not arise independently of the question of what goals it is intended to serve. We cannot fully understand a theoretical perspective until we understand the context in which it is proposed.

This completes the case for perspectivism.

References

- Arthur, R. W. T. (2016). "On the mathematization of free fall: Galileo, Descartes, and the history of a misconstrual." In *The Language of Nature: Reassessing the Mathematization of Natural Philosophy in the Seventeenth Century*, eds. G. Gorham, B. Hill, E. Slowik, and C. K. Waters: 81-111. Minneapolis: University of Minnesota Press.
- Bogen, J. and Woodward, J. (1988). "Saving the phenomena." *Philosophical Review* 97(3): 303-352.
- Boyd, R. (1983). "Realism, approximate truth, and scientific method." In *Scientific Theories*, ed. C. W. Savage: 355-391. Minneapolis: University of Minnesota Press.
- Boyd, R. (1999). "Homeostasis, species, and higher taxa." In *Species: New Interdisciplinary Essays*, ed. R. A. Wilson: 141-185. Cambridge: The MIT Press.
- Boyd, R. (2010). "Homeostasis, higher taxa, and monophyly." *Philosophy of Science* 77(5): 686-701.
- Boyd, R. (2019). "Rethinking natural kinds, reference, and truth." *Synthese* 198 (suppl.12): 2863-2903.
- Bradbury, S. (1967). *The Evolution of the Microscope*. Oxford: Pergamon Press.
- Brandon, R. (2010) "Adaptation and evolutionary theory." In *Philosophy of Biology: An Anthology*, eds. A. Rosenberg and R. Arp: 103-120. Oxford: Wiley-Blackwell.
- Bursten, J. R. (2018). "Smaller than a breadbox: Scale and natural kinds." *British Journal for Philosophy of Science* 69(1): 1-23.
- Button, T. (2013). *The Limits of Realism*. Oxford: Oxford University Press.

Byrne, A. and Hilbert, D. R. (2003). "Color realism and color science." *Behavioral and Brain Sciences* 26(1): 3-21.

Cadenasso, M. L., Pickett, S. T. A., Weathers, K. C., and Jones, C. G. (2003). "A framework for a theory of ecological boundaries." *BioScience* 53(8): 750-758.

Cartwright, N. (1983). *How the Laws of Physics Lie*. Oxford: Oxford University Press.

Cartwright, N. (1999). *The Dappled World: A Study of the Boundaries of Science*. Cambridge: Cambridge University Press.

Chakravartty, A. (2007). *A Metaphysics for Scientific Realism*. Cambridge: Cambridge University Press.

Chakravartty, A. (2010). "Perspectivism, inconsistent models, and contrastive explanation." *Studies in the History and Philosophy of Science Part A* 41(4): 405-412.

Chakravartty, A. (2011). "The semantic or model-theoretic view of theories and scientific realism." *Synthese* 127(3): 325-345.

Chakravartty, A. (2017a). "Scientific realism." In *The Stanford Encyclopedia of Philosophy*, ed. E. Zalta. <https://plato.stanford.edu/entries/scientific-realism/>

Chakravartty, A. (2017b). *Scientific Ontology: Integrating Naturalized Metaphysics and Voluntarist Epistemology*. Oxford: Oxford University Press.

Chang, H. (2004). *Inventing Temperature: Measurement and Scientific Progress*. Oxford: Oxford University Press.

Chang, H. (2009). "Ontological principles and the intelligibility of epistemic activities." In *Scientific Understanding: Philosophical Perspectives*, eds. H. de Regt, S. Leonelli, and K. Eigner: 64-82. Pittsburgh, PA: University of Pittsburgh Press.

Chang, H. (2017). "Operational coherence as the source of truth." *Proceedings of the Aristotelian Society* 117(2): 103-122.

Chang, H. (2018). "Realism for realistic people." *Spontaneous Generations: A Journal for the History and Philosophy of Science* 9(1): 31-34.

Chirimuuta, M. (2016). "Vision, perspectivism, and haptic realism." *Philosophy of Science* 83(5): 746-756.

Clark, A. and Chalmers, D. J. (1998). "The extended mind." *Analysis* 58(1): 7-19.

Clark, A. (2003). *Natural Born Cyborgs*. Oxford: Oxford University Press.

Clark, A. (2008). *Supersizing the Mind*. Oxford: Oxford University Press.

Clark, A. (2010). "Coupling, constitution, and the cognitive kind: A reply to Adams and Aizawa." In *The Extended Mind*, ed. R. Menary: 81-100. Cambridge, MA: MIT Press.

Clark, A. & Chalmers, D. (1998). "The extended mind." *Analysis* 58(1): 7-19.

Collins, H. (2010). "Humans not instruments." *Spontaneous Generations: A Journal for the History and Philosophy of Science* 4(1): 138-147.

Contessa, G. (2010). "Empiricist structuralism, metaphysical realism, and the bridging problem." *Analysis* 70(3): 514-524.

Cretu, A. M. (2018). "Perspectival realism." In *Encyclopedia of Educational Philosophy and Theory*, ed. M. A. Peters. Singapore: Springer.

Cretu, A. M. (2020). "Natural kinds as real patterns: Or how to solve the commitment problem for perspectival realism." [preprint]

Danks, D. (2019). "Safe and substantive perspectivism." In *Understanding Perspectivism: Scientific Challenges and Methodological Prospects*, eds. M. Massimi and C. D. McCoy: 127-140. New York: Routledge.

Dear, P. (2006). *The Intelligibility of Nature: How Science Makes Sense of the World*. Chicago: University of Chicago Press.

Dennett, D. C. (1991). "Real patterns." *Journal of Philosophy* 88(1): 27-51.

Dick, S. J. (2013). *Discovery and Classification in Astronomy: Controversy and Consensus*. Cambridge: Cambridge University Press.

Dimech, D. K. (2017). "Modeling with words: Narrative and natural selection." *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biology and Biomedical Sciences* 62:20-24.

Dupre, J. (1993). *The Disorder of Things: Metaphysical Foundations of the Disunity of Science*. Cambridge, MA: Harvard University Press.

Elgin, C. (2019) "Nominalism, realism, and objectivity." *Synthese* 196(2): 519-534.

Ellis, B. D. (2001). *Scientific Essentialism*. Cambridge: Cambridge University Press.

Ereshefsky, M. (1992). "Eliminative pluralism." *Philosophy of Science* 59(4): 671-690.

Fellowes, S. (2020). "Scientific perspectivism and psychiatric diagnoses: Respecting history and constraining relativism." *European Journal for Philosophy of Science* 11(1): 1-24.

Field, H. (1980). *Science Without Numbers*. Oxford: Oxford University Press.

Fine, A. (1996) *The Shaky Game: Einstein, Realism, and the Quantum Theory*. Chicago: University of Chicago Press.

Forinash, K., Rumsey, W., and Lang, C. (2000) "Galileo's mathematical language of nature." *Science & Education* 9: 449-457.

Franklin-Hall, L. (2015). "Natural kinds as categorical bottlenecks." *Philosophical Studies* 172(4): 925-948.

Freedman, W. J. and Kaufmann III, R. A. (2008). *Universe*, 8th edition. London: London Bridge Books.

Frigg, R. (2006). "Scientific representation and the semantic view of theories." *Theoria* 21(1): 49-65.

Frigg, R. and Votsis, I. (2011) "Everything you always wanted to know about structural realism but were afraid to ask." *European Journal for Philosophy of Science* 1(2): 227-276.

Gannett, L. (2003) "Making populations: Bounding genes in space and time." *Philosophy of Science* 70: 989-1001.

Garber, D. (2016) "Laws of nature and the mathematics of motion." In *The Language of Nature: Reassessing the Mathematization of Natural Philosophy in the Seventeenth Century*, eds. G. Gorham, B. Hill, E. Slowik, and C. K. Waters: 134-159. Minneapolis: University of Minnesota Press.

Gibson, J. J. (1979). *The Ecological Approach to Visual Perception*. London: Houghton Mifflin.

Giere, R. N. (1988). *Explaining Science: A Cognitive Approach*. Chicago: University of Chicago Press.

Giere, R. N. (1999). "Naturalism and realism." In *Science Without Laws*, ed. R. N. Giere: 69-83. Chicago and London: The University of Chicago Press.

Giere, R. N. (2005). "Scientific realism: Old and new problems." *Erkenntnis* 63: 149-165.

Giere, R. N. (2006). *Scientific Perspectivism*. Chicago and London: The University of Chicago Press.

Giere, R. N. (2009). "Scientific perspectivism: Behind the stage door." *Studies in History and Philosophy of Science Part A* 40(2): 221-223.

Giere, R. N. (2010). "An agent-based conception of models and scientific representation." *Synthese* 172: 269–281.

Giere, R. N. (2012). "Kuhn as perspectival realist." *Topoi* 32(1): 53-57.

Giere, R. N. (2015). "Contingency, conditional realism, and the evolution of the sciences." In *Science as it Could Have Been: Discussing the Contingency/Inevitability Problem*, eds. L. Soler, E. Trizio, and A. Pickering: 187-201. Pittsburgh: University of Pittsburgh Press.

Giere, R. N. (2016) "Feyerabend's perspectivism." *Studies in History and Philosophy of Science Part A* 57: 137-141.

Godfrey-Smith, P. (2009). "Abstractions, idealizations, and evolutionary biology." In *Mapping the Future of Biology: Evolving Concepts and Theories*, eds. A. Barberousse, M. Morange, and T. Pradeu: 47-56. Dordrecht: Springer.

Goodman, N. (1972). "Seven strictures on similarity." In *Problems and Projects*, ed. N. Goodman: 437-446. Indianapolis: Bobbs-Merrill.

Goodman, N. (1978). *Ways of Worldmaking*. Hassocks: Harvester Press.

Grant, P. R. (1986). *Ecology and Evolution of Darwin's Finches*. Princeton, N.J.: Princeton University Press.

Hacking, I. (1983). *Representing and Intervening*. Cambridge: Cambridge University Press.

Hacking, I. (1999). *The Social Construction of What?* Cambridge, MA: Harvard University Press.

Hand, D. J. (2004). *Measurement Theory and Practice: The World Through Quantification*. London: Arnold.

Hardin, C. L. (1988). *Color for Philosophers*. Indianapolis: Hackett.

Harré, R. (2003). "The materiality of instruments in a metaphysics for experiments." In *The Philosophy of Scientific Experimentation*, ed. H. Radder: 19-38. Pittsburgh: The University of Pittsburgh Press.

Harré, R. (2010). "Equipment for an experiment." *Spontaneous Generations: A Journal for the History and Philosophy of Science* 4(1): 30-38.

Hellman, G. (1989). *Mathematics Without Numbers: Towards a Modal-Structural Approach*. Oxford: Clarendon Press.

Hirsch, E. (2005). "Physical-object ontology, verbal disputes, and common sense." *Philosophy and Phenomenological Research* 70(1): 67-97.

Hofer, C. (2019) "Scientific realism without the quantum." In *Scientific Realism and the Quantum*, eds. S. French and J. Saatsi: 19-34. Oxford: Oxford University Press.

Hoehn, R. D., Nichols, D. E., Neven, H., and Kais, S. (2018). "Status of the vibrational theory of olfaction." *Frontiers in Physics* 6: 1-25.

Hudson, R. G. (1994). "Reliability, pragmatic and epistemic." *Erkenntnis* 40(1): 71-86.

Humphreys, P. (2004). *Extending Ourselves: Computational Science, Empiricism, and Scientific Method*. New York: Oxford University Press.

Humphreys, P. (2013). "Scientific ontology and speculative ontology." In *Scientific Metaphysics*, eds. D. Ross, J. Ladyman, and H. Kincaid: 51-78. Oxford: Oxford University Press.

Jesseph, D. (2016). "Ratios, quotients, and the language of nature." In *The Language of Nature: Reassessing the Mathematization of Natural Philosophy in the Seventeenth Century*, eds. G. Gorham, B. Hill, E. Slowik, and C. K. Waters: 160-177. Minneapolis: University of Minnesota Press.

Kabay, P. (2010). *On the Plenitude of Truth: A Defense of Trivialism*. LAP Lambert Academic Publishing.

Kay, P. and Maffi, L. (1999). "Color appearance and the emergence and evolution of basic color lexicons." *American Anthropologist* 101(4): 743-760.

Kirkham, R. L. (1992). *Theories of Truth: A Critical Introduction*. Cambridge, MA: The MIT Press.

Kitcher, P. (1989). "Explanatory unification and the causal structure of the world." In *Scientific Explanation*, eds. P. Kitcher and W. Salmon: 410-505. Minneapolis: University of Minnesota Press.

Kitcher, P. (1993). *The Advancement of Science: Science Without Legend, Objectivity Without Illusions*. Oxford: Oxford University Press.

Kitcher, P. (2001). "Real realism: The Galilean strategy." *The Philosophical Review* 110(2): 151-197.

Kitcher, P. (2002). "On the explanatory role of correspondence truth." *Philosophy and Phenomenological Research* 64(2): 346-364.

Koslicki, K. (2008). "Natural kinds and natural kind terms." *Philosophy Compass* 3(4): 789-802.

Kosso, P. (2002). "The omniscienter: Beauty and scientific understanding." *International Studies in the Philosophy of Science* 16(1): 39-48.

Kuhn, T. S. (2000) *The Road Since Structure: Philosophical Essays, 1970-1993*, ed. J. Conant and J. Haugeland. Chicago: University of Chicago Press.

Kukla, A. (2000). *Social Constructivism and the Philosophy of Science*. New York: Routledge.

Kusch, M. (2017). "Epistemic relativism, scepticism, pluralism." *Synthese* 194(12): 4687-4703.

Lack, D. (1945). *The Galapagos Finches (Geospizinae): A Study in Variation*. San Francisco: California Academy of Sciences.

Ladyman, J. and Ross, D. (2007). *Every Thing Must Go: Metaphysics Naturalized*. Oxford: Oxford University Press.

Laudan, L. (1977). *Progress and Its Problems: Towards a Theory of Scientific Growth*. Berkeley: University of California Press.

Laudan, L. (1981). "A confutation of convergent realism." *Philosophy of Science* 48(1): 19-49.

Levins, R. (1966). "The strategy of model building in population biology." *American Scientist* 54: 421-431.

Lewontin, R. C. (1970). "The units of selection." *Annual Reviews of Ecology and Systematics* 1: 1-18.

Lidgard, S. and Nyhart, L. K. (2017). "The work of biological individuality: Concepts and contexts." In *Biological Individuality: Integrating Scientific, Philosophical, and Historical Perspectives*, eds. S. Lidgard and K. L. Nyhart: 17-62. Chicago: The University of Chicago Press.

Lipton, P. (2002). "The reach of the law." *Philosophical Books* 43(4): 254-260.

Lyon, A. (2014). "Data." In *The Oxford Handbook of Philosophy of Science*, ed. P. Humphreys: 738-759. New York, NY: Oxford University Press.

Massimi, M. (2012). "Scientific perspectivism and its foes." *Philosophica* 84: 25-52.

Massimi, M. (2016a). "Four kinds of perspectival truth." *Philosophy and Phenomenological Research* 96(2): 342-359.

Massimi, M. (2016b). "Bringing real realism back home: A perspectival slant." In *The Philosophy of Philip Kitcher*, eds. M. Couch and J. Pfeifer: 98-120. New York: Oxford University Press.

Massimi, M. (2017). "Perspectivism." In *The Routledge Handbook of Scientific Realism*, ed. J. Saatsi: 164-175. New York: Routledge.

Massimi, M. and C. D. McCoy, eds. (2019). *Understanding Perspectivism: Scientific Challenges and Methodological Prospects*. New York: Routledge.

Masterman, M. (1970). "The nature of a paradigm." In *Criticism and the Growth of Knowledge*, eds. I. Lakatos and A. Musgrave: 59-90. Cambridge: Cambridge University Press.

Maxwell, G. (1962). "The ontological status of theoretical entities." In *Scientific Explanation, Space, and Time*, eds. H. Fiegl and G. Maxwell: 181-192. Minneapolis: University of Minnesota Press.

McAllister, J. W. (1997). "Phenomena and patterns in data sets." *Erkenntnis* 47(2): 217-228.

McCloskey, M. (1983). "Naive theories of motion." In *Mental Models*, eds. D. Getner and A. L. Stevens: 299-324. Hillsdale: Erlbaum.

McGrath, M. and Frank, D. (2018). "Propositions." In *The Stanford Encyclopedia of Philosophy*, ed. E. Zalta. <https://plato.stanford.edu/entries/propositions/>

McMullin, E. (1985). "Galilean Idealization." *Studies in History and Philosophy of Science Part A* 16(3): 247-273.

McMullin, E. (1993). "Rationality and paradigm change in science." In *World Changes: Thomas Kuhn and the Nature of Science*, ed. P. Horwich: 55-78. Cambridge, MA: The MIT Press.

Morgan, M. S., and Morrison, M. (Eds.). (1999). *Models as Mediators: Perspectives on Natural and Social Science*. Cambridge: Cambridge University Press.

Morrison, M. (2011). "One phenomenon, many models." *Studies in History and Philosophy of Science Part A* 42(2): 342-351.

Morrison, M. (2015). *Reconstructing Reality: Models, Mathematics, and Simulations*. New York: Oxford University Press.

Muller, F. A. and van Fraassen, B. C. (2008). "How to talk about unobservables." *Analysis* 68(3): 197-205.

Musgrave, A. (1989). "NOA's ark – Fine for realism." *Philosophical Quarterly* 39(157): 383-398.

Nagel, T. (1986). *The View from Nowhere*. Oxford: Oxford University Press.

Nguyen, J. (2020). "It's not a game: Accurate representation and toy models." *British Journal for the Philosophy of Science* 71(3): 1013-1041.

Nguyen, J. and Frigg, R. (2017). "Mathematics is not the only language in the book of nature." *Synthese*: 1-22.

Niiniluoto, I. (1999). *Critical Scientific Realism*. Oxford: Oxford University Press.

Noë, A. (2004). *Action in Perception*. Cambridge, MA: MIT Press.

Oddie, G. (2014). "Truthlikeness." In *The Stanford Encyclopedia of Philosophy*, ed. E. Zalta. <https://plato.stanford.edu/entries/truthlikeness/>

Odenbaugh, J. (2011). "True lies: Realism, robustness, and models." *Philosophy of Science* 78(5): 1177-1188.

Ordorica, S. A. (2016). "The explanatory role of abstraction processes in models: The case of aggregations." *Studies in History and Philosophy of Science Part A* 56: 161-167.

Oreskes, N. (2007). "From scaling to simulation: Changing meanings and ambitions of models in geology." In *Science Without Laws: Model Systems, Cases, Exemplary Narratives*, eds. A. N. H. Creager, E. Lunbeck, and M. N. Wise: 93-124. Durham and London: Duke University Press.

Pinch, T. (1985). "Towards an analysis of scientific observation: The externality and evidential significance of observation reports in physics." *Social Studies of Science* 15(1): 3-36.

Popper, K. R. (1960). *The Poverty of Historicism*. London: Routledge & Kegan Paul.

Psillos, S. (1999). *Scientific Realism: How Science Tracks Truth*. London: Routledge.

Psillos, S. (2011). "Living with the abstract: Realism and models." *Synthese* 180(1): 3-17.

Putnam, H. (1996). "Is there anything still to say about reality and truth?" In *Starmaking: Realism, Anti-Realism, and Irrationalism*, ed. P. J. McCormick: 11-28. Cambridge, MA: The MIT Press.

Quine, W. V. O. (1960). *Word and Object*. Cambridge, MA: The MIT Press.

Reiss, J. and Sprenger, J. (2020). "Scientific objectivity." In *The Stanford Encyclopedia of Philosophy*, ed. E. Zalta. <https://plato.stanford.edu/entries/scientific-objectivity/>

Reynolds, A. S. (2018). *The Third Lens: Metaphor and the Creation of Modern Cell Biology*. Chicago: University of Chicago Press.

Ross, P. W. (2015). "Primary and secondary qualities." In *The Oxford Handbook of the Philosophy of Perception*, ed. M. Matthen: 405-412. Oxford: Oxford University Press.

Rothbart, D. (2007). *Philosophical Instruments: Minds and Tools at Work*. Urbana & Chicago: The University of Illinois Press.

Rueger, A. (2016). "Perspectival realism and incompatible models." *Axiomathes* 26(4): 401-410.

Ruphy, S. (2016). *Scientific Pluralism Reconsidered: A New Approach to the Disunity of Science*. Pittsburgh: University of Pittsburgh Press.

Ruse, M. (1992). "Biological species: Natural kinds, individuals, or what?" In *The Units of Evolution: Essays on the Nature of Species*, ed. M. Ereshefsky: 343-362. Cambridge: The MIT Press.

Russell, B. (1912). *The Problems of Philosophy*. Oxford: Oxford University Press.

Russell, B. (1927). *The Analysis of Matter*. London: George Allen & Unwin.

Ruyant, Q. (2020). "Persectival realism and norms of scientific representation." *European Journal for Philosophy of Science* 10(20): 1-17.

Sankey, H. (2008). "Scientific realism and the inevitability of science." *Studies in the History and Philosophy of Science Part A* 39(2): 259-264.

Schindler, S. (2011). "Bogen and Woodward's data-phenomena distinction, forms of theory-ladenness, and the reliability of data." *Synthese* 182(1): 39-55.

Shaffer, M. J. (2012). *Counterfactuals and Scientific Realism*. New York: Palgrave Macmillan.

Sismondo, S. and Chrisman, N. (2001). "Deflationary metaphysics and the natures of maps." *Philosophy of Science* 68(3): S38-S49.

Smith, B. and Casati, R. (1994). Naive physics: An essay in ontology. *Philosophical Psychology* 7(2): 227-247.

Sober, E. (1980). "Evolution, population thinking, and essentialism." *Philosophy of Science* 47(3): 350-383.

Stanford, P. K. (2006). *Exceeding Our Grasp: Science, History, and the Problem of Unconceived Alternatives*. Oxford: Oxford University Press.

Stanford, P. K. (2017). "So long, and thanks for all the fish: Metaphysics and the philosophy of science." In *Metaphysics and the Philosophy of Science: New Essays*, eds. M. Slater and Z. Yudell: 127-140. Oxford: Oxford University Press.

Stevens, S. S. (1946). "On the theory of scales of measurement." *Science* 103: 677-680.

Strayer, D. L., Power, M. E., Fagan, W. F., Pickett, S. T. A., and Belnap, J. (2003). "A classification of ecological boundaries." *BioScience* 53(8): 723-729.

Suppes, P. (1962). "Models of data." In *Logic, Methodology, and Philosophy of Science: Proceedings of the 1960 International Congress*, eds. E. Nagel, P. Suppes, and A. Tarski: 252-261. Stanford: Stanford University Press.

Sylvan, R. (2003). "The importance of nonexistent objects and of intensionality in mathematics." *Philosophia Mathematica* 11(1): 20-52.

Tal, E. (2013). "Old and new problems in philosophy of measurement." *Philosophy Compass* 8(12): 1159-1173.

Taylor, H. (2020). "Emotions, concepts, and the indeterminacy of natural kinds." *Synthese* 197(5): 2073-2093.

Taylor, H. (2021). "Whales, fish, and Alaskan bears: Interest-relative taxonomy and kind pluralism in biology." *Synthese* 198: 3369-3387.

Teller, P. (2016). "Pan-perspectival realism explained and defended"

Teller, P. (2018a). "Referential and perspectival realism." *Spontaneous Generations: A Journal for the History and Philosophy of Science* 9(1): 151-164.

Teller, P. (2018b). "Measurement accuracy realism." In *The Experimental Side of Modeling*, eds. I. F. Peschard and B. C. van Fraassen: 273-298. Minneapolis: University of Minnesota Press.

Teller, P. (2019). "What is perspectivism, and does it count as realism?" In *Understanding Perspectivism: Scientific Challenges and Methodological Prospects*, eds. M. Massimi and C. D. McCoy: 49-64. New York: Routledge.

Thomson-Jones, M. (2010). "Missing systems and the face-value practice." *Synthese* 172(2): 283-299.

Thomson-Jones, M. (2012). "Modeling without mathematics." *Philosophy of Science* 79(5): 761-772.

Tinsley, J. N., Molodstov, M. I., Prevedel, R., Wartmann, D., Espigule-Pons, J., Lauwers, M., and Vaziri, A. (2016). "Direct detection of a single photon by humans." *Nature Communications* 7(12172): 1-25.

Toon, A. (2014). "Empiricism for cyborgs." *Philosophical Issues* 24(1): 409-425.

Tulodziecki, D. (2016). "Structural realism beyond physics." *Studies in History and Philosophy of Science Part A* 59: 106-114.

Turner, D. D. (2007). *Making Prehistory: Historical Science and the Scientific Realism Debate*. Cambridge: Cambridge University Press.

van Fraassen, B. C. (1980). *The Scientific Image*. Oxford: Oxford University Press.

van Fraassen, B. C. (2001). "Constructive empiricism now." *Philosophical Studies* 106(1): 151-170.

van Fraassen, B. C. (2008). *Scientific Representation: Paradoxes of Perspective*. Oxford: Oxford University Press.

Vertesi, J. (2015). *Seeing Like A Rover: How Robots, Teams, and Images Craft Knowledge of Mars*. Chicago: The University of Chicago Press.

Votsis, I. (2012). "Putting realism in perspectivism." *Philosophica* 84: 85-122.

Walton, K. (1984). "Transparent pictures: On the nature of photographic realism." *Critical Inquiry* 11(2): 246-277.

Weiner, J. (1995). *The Beak of the Finch: A Story of Evolution in Our Time*. New York: Knopf.

Weisberg, M. (2006). "Water is not H₂O". In *Philosophy of Chemistry: Synthesis of a New Discipline*, eds. D. Baird, E. Scerri, and L. McIntyre: 337-345. Dordrecht: Springer.

Weisberg, M. (2007). "Three kinds of idealization." *The Journal of Philosophy* 104(12): 639-659.

Whewell, W. (1860). *On the Philosophy of Discovery*. London: Chapters Historical and Critical.

Wimsatt, W. (2007). *Re-Engineering Philosophy for Limited Beings*. Cambridge, Mass.: Harvard University Press.

Woolgar, S. (1988). *Science: The Very Idea*. London: Tavistock.

Worrall, J. (1989). "Structural realism: The best of both worlds?" *Dialectica* 43(1-2): 99-124.

Wray, K. B. (2015). "Pessimistic induction: Four varieties." *International Studies in the Philosophy of Science* 29(1): 61-73.

Wray, K. B. (2018). *Resisting Scientific Realism*. Cambridge: Cambridge University Press.