



THE UNIVERSITY *of* EDINBURGH

This thesis has been submitted in fulfilment of the requirements for a postgraduate degree (e.g. PhD, MPhil, DClinPsychol) at the University of Edinburgh. Please note the following terms and conditions of use:

- This work is protected by copyright and other intellectual property rights, which are retained by the thesis author, unless otherwise stated.
- A copy can be downloaded for personal non-commercial research or study, without prior permission or charge.
- This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the author.
- The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the author.
- When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given.

A Nominalist's Credo

James Henry Collin

For Research Degree of Ph.D. in Philosophy

University of Edinburgh

2012

Abstract

Introduction: I lay out the broad contours of my thesis: a defence of mathematical nominalism, and nominalism more generally. I discuss the possibility of metaphysics, and the relationship of nominalism to naturalism and pragmatism.

Chapter 2: I delineate an account of abstractness. I then provide counter-arguments to claims that mathematical objects make a difference to the concrete world, and claim that mathematical objects are abstract in the sense delineated.

Chapter 3: I argue that the epistemological problem with abstract objects is not best understood as an incompatibility with a causal theory of knowledge, or as an inability to explain the reliability of our mathematical beliefs, but resides in the epistemic luck that would infect any belief about abstract objects. To this end, I develop an account of epistemic luck that can account for cases of belief in necessary truths and apply it to the mathematical case.

Chapter 4: I consider objections, based on (meta)metaphysical considerations and linguistic data, to the view that the existential quantifier expresses existence. I argue that these considerations can be accommodated by an existentially committing quantifier when the pragmatics of quantified sentences are properly understood. I develop a semi-formal framework within which we can define a notion of nominalistic adequacy. I show how our notion of nominalistic adequacy can show why it is legitimate for the nominalist to make use of platonistic “assumptions” in inference-making.

Chapter 5: I turn to the application of mathematics in science, including explanatory applications, and its relation to a number of indispensability arguments. I consider also issues of realism and anti-realism, and their relation to these arguments. I argue that abstraction away from pragmatic considerations has acted to skew the debate, and has obscured possibilities for a nominalistic understanding of mathematical practices. I end by explaining the notion of a pragmatic meta-vocabulary, and argue that this notion can be used to carve out a new way of locating our ontological commitments.

Chapter 6: I show how the apparatus developed in earlier chapters can be utilised to roll out the nominalist project to other domains of discourse. In particular, I consider propositions and types. I claim that a unified account of nominalism across these domains is available.

Conclusion: I recapitulate the claims of my thesis. I suggest that the goal of mathematical enquiry is not descriptive knowledge, but understanding.

Without a continual falsification of the world by means of numbers, mankind could not live.

- Friedrich Nietzsche, *Beyond Good and Evil*

Clearly human beings could dispense with all discourse, though only at the expense of having nothing to say.

- Wilfred Sellars, 'A Semantical Solution to the Mind-Body Problem'.

Contents

Abstract	2
Acknowledgements	7
1 Introduction	8
1.1 Profanity and the Possibility of Metaphysics	8
1.2 The Credo	11
1.3 Nominalism and Pragmatism	12
1.4 Outlining the Project	14
2 Abstract Objects	16
2.1 Mathematical Objects and Causation	16
2.2 Mathematical Objects and Conservativeness	19
2.3 Mathematical Objects and Laws	20
3 Epistemic Luck and Necessary Truths	25
3.1 Prologue	25
3.2 The Luck Platitude	27
3.2.1 Safety	28
3.2.2 The Problem	28
3.3 Epistemic Possibility	31
3.3.1 Refinements	33
3.4 Putting Epistemic Safety to Work	36
3.4.1 Irrelevance	36
3.4.2 Mathematical Objects	37
4 Existential Quantification: The Third Way	42
4.1 The Quinian Orthodoxy	42
4.2 Doubting the Quinian Orthodoxy	43
4.2.1 Metametaphysics (I)	43

<i>CONTENTS</i>	5
4.2.2 Metametaphysics (II)	44
4.2.3 Linguistic Data	45
4.3 All Roads Lead to Quine	47
4.3.1 Trivializing Truth-Conditions	47
4.3.2 A Quantifier for Metaphysics	48
4.4 The Third Way	49
4.4.1 Two Ways of Thinking About Communication	50
4.4.2 Embellished Worlds	53
4.5 Applications	56
4.5.1 Losing Our Grip on the World	59
4.6 Summation	60
5 Mathematics and the World	61
5.1 The Indispensability Argument	61
5.1.1 Premiss 4	62
5.1.2 Premiss 2	63
5.1.3 Premiss 1	66
5.1.4 Premiss 3	73
5.2 Indispensability and Pragmatism	79
5.2.1 The Enhanced Indispensability Argument	83
5.3 Applying Mathematics	87
5.3.1 How Scientists Represent the World	87
5.3.2 The Possibility of Applied Mathematics	90
5.3.3 Theories of Applied Mathematics in the Literature	93
5.3.4 The Inferential Account	95
5.3.5 Issues	98
5.3.6 How Modelling Provides Epistemic Access to the World	109
5.4 What is Explanation?	114
5.4.1 Prominent Models of Explanation	115
5.4.2 Explanation and Understanding	118
5.4.3 Understanding and Explanation	126
5.4.4 The Enhanced Indispensability Argument Revisited	128
5.5 Renewing Ontology	130
5.5.1 Pragmatic Meta-Analysis	131
6 Rolling Out the Project	135
6.1 Propositions as Tools of Measurement	136
6.1.1 Developing the Measurement-Theoretic Account	139

6.1.2	Propositions and the Pragmatics of Representation	142
6.1.3	Must Propositional Attitudes be Relations?	144
6.1.4	Are Measurement-Theoretic Accounts Inevitable?	146
6.2	Types and Tokens	147
6.2.1	Instances of Type-Talk	148
6.2.2	Difficulties with Paraphrase Strategies	149
6.2.3	Problems for Type Platonism	151
6.2.4	What is a Type?	155
6.2.5	Type Nominalism	157
7	Conclusion	161
7.1	An Autobiographical Note	161
7.2	A Nominalist's Credo	162

Acknowledgements

Many people amongst the postgraduate group and philosophy faculty at Edinburgh are owed a debt of gratitude for their assistance, feedback, and conversation. In particular I must thank Jesper Kallestrup for his supervision and encouragement and William Stirton for stepping in late on in the process, and for the thoughtful and perspicacious commentary he provided. Particular thanks are owed to Jeff Ketland. No-one has done more to shape my own philosophical outlook, even if some of its lineaments have been a continuing source of dismay to him. In his instruction, guidance and counsel, his generosity of spirit always outweighed his (not inconsiderable) dislike of nominalism. A different kind of gratitude is owed to my wife Ailsa, for her unerring love and support during my time as a graduate student. Your grace made this possible.

Chapter 1

Introduction

Nominalists believe that there are no abstract objects. One important class of abstract objects are mathematical objects—numbers, sets, functions and so on—but there are other important kinds of abstracta besides, such as propositions, and types that also play a role in philosophizing. The nominalist believes that there are fewer things in heaven and earth than are dreamt of in our philosophy. Claims such as these are very much the province of philosophy, but what can one find out about the world from the armchair? We philosophers who engage with ontological questions take our place in part of a venerable tradition of arguing about what the world is like, but it may, and perhaps should, seem odd that philosophers would think to engage in such a practice at all. When, to take but one example, a new species of beetle is discovered, this does not happen a priori; *Anthonomus grandis* had to be carefully observed and documented before any entomologist worth the name thought to include boll weevils in her ontology. How is it then that what philosophers do could hope to tell us what kind of a world we live in? Metaphysics is the philosophical study of reality, so if we are to provide some vindication of the viability of metaphysics we must consider what it is that makes a practice a peculiarly philosophical one, and ask why it is that a thing like that could be used to reveal something about the world. While it is par for the course to define one's terms at the beginning of a philosophical monograph, defining philosophy itself might strike the reader as a tad gratuitous, but here I think it is informative, particularly in light of an ambient philosophical self-doubt about the possibility of metaphysics impelled sometimes by a lingering positivism, sometimes by a naturalistic aversion to “scholastic” metaphysics.

1.1 Profanity and the Possibility of Metaphysics

In a well known passage, Michael Dummett discusses the derogatory racial term “boche”; its *circumstances* of application, its *consequences* of application, and their relationship to its meaning. While Dummett's interests here are specifically in the philosophy of language, his observations will provide a

good route into the aims and methods of philosophy more generally:

A simple case would be that of a pejorative term, e.g. ‘Boche’. The condition for applying the term to someone is that he is of German nationality; the consequences of its application are that he is barbarous and more prone to cruelty than other Europeans. We should envisage the connections in both directions as sufficiently tight as to be involved in the very meaning of the word: neither could be severed without altering its meaning. Someone who rejects the word does so because he does not want to permit a transition from the grounds for applying the term to the consequences of doing so. The addition of the term ‘Boche’ to a language which did not previously contain it would produce a non-conservative extension, i.e. one in which certain other statements which did not contain the term were inferable from other statements not containing it which were not previously inferable. [Dummett 1973: 454]

In a commentary on the passage, Robert Brandom makes a number of points pertinent to our current concerns. In the first place, concepts can be criticized on the basis of substantive beliefs. If one knows of Germans who are not unusually barbarous and cruel, then one has reason to take issue with the concept Boche, but *how* one must take issue with the concept is informative. One cannot claim that there are no Boche, for the circumstances of applying the term to a person are when he or she is German, and hence denying that there are Boche is just to deny that there are Germans. Nor, on either pain of incoherence or simply changing the meaning of the term, can one allow that there are Boche, but deny that they are cruel, for the conceptual consequence of being a Boche just is that one is cruel. One must renounce the concept Boche itself; refuse to employ it on the grounds that it manifests an inference that is defective, that should not be endorsed. Evaluating ‘Boche’ involves making explicit the inferential commitments that had hitherto lain implicit in the concept. Quite generally, by excavating the commitments of a concept in this way and laying them bare, one brings it into the light, and makes possible the evaluation of the concept itself. In this way it can be subject to demands for justification and defence, and may or may not be found wanting.¹ This I take it is a model of philosophy. This process of standing back from thought and bringing it under scrutiny—what Putnam [2004] has called ‘reflective transcendence’—transforms that which *drives* one’s thinking into something that can itself be thought about. G.K. Chesterton put it succinctly when he said:

Philosophy is merely thought that has been thought out. It is often a great bore. But man has no alternative, except between being influenced by thought that has been thought out and being influenced by thought that has not been thought out. [Chesterton 1950: 180]

The philosopher’s task is to clarify our commitments, and, in doing so, to see how things fit together, to locate our beliefs in a wider network of inferential relations. Philosophy is aimed at acquiring a certain kind of *understanding*, an integration of our beliefs into a coherent whole.² This then is what philosophy

¹Brandom adds, rather inspiringly: ‘In Reason’s fight against thought debased by prejudice and propaganda, the first rule is that material inferential commitments that are potentially controversial should be made explicit as claims, exposing them both as vulnerable to reasoned challenge and as in need of reasoned defense.’ [Brandom 1994: 126]

²This provides something of a vindication of the analytic philosopher’s sometimes criticised twin preoccupations with language and with logic; for it is through language that we stake out our commitments and their prerequisites and entailments, and it is in formal methods that the clarity and precision with which we make these things explicit achieves its apotheosis.

can tell us about reality: careful philosophical work can excavate our commitments to the world being a certain way, to the existence of objects or kinds of objects; commitments that we were previously unaware of. More strongly, it can show us that we are warranted in taking the world to be a certain way, or in taking certain kinds of objects to exist. We can understand the battle lines between the platonist and the nominalist to be drawn thus: platonists hold either that we are committed to the existence of abstract objects in virtue of other things that we accept, or ought to accept, or that we are warranted in taking abstract objects to exist in virtue of other things that we accept or ought to accept. Nominalists deny this.

Although the preceding gives us some idea of how metaphysics is *possible*, it also gives us some idea of why philosophers, notoriously, fail to achieve any sort of consensus on metaphysical, or other, matters. Unearthing our inferential commitments can only take us so far; it can tell us that ‘p’, ‘ $p \rightarrow q$ ’ and ‘ $\neg q$ ’ are not compossible, but it does not tell us what thereby to *do*; whether to reject our belief that p, our belief that $\neg q$, or to reconsider the conditional itself. As the saying goes, one man’s *modus ponens* is another man’s *modus tollens*. For this reason we should expect progress in philosophy to consist in an ever more finely grained understanding of how some commitments and positions are related to others, and in this philosophers indeed make progress. Both the thesis that philosophy does not progress because it does not achieve consensus or resolution, and the objection to philosophy as a practice on those grounds, subsequently miss the mark. This means that I do not hold out much hope that platonists who might read this thesis will be converted, not because platonists are unusually dogmatic, but only because it is not ordinarily within the powers or the remit of mere philosophy to do such a thing. I do hope however to excavate some of the fundamental assumptions that drive this debate, and, by exposing them to scrutiny, show that much of that which is taken to undermine nominalism lacks the force it may appear, on first inspection, to possess.

Talk of naturalism abounds in the debate between nominalists and realists, just as it does elsewhere in philosophy. My aim here is not to defend a naturalistic account of nominalism *per se*, but from what we have said about the character of philosophy, there is a sense in which naturalism *of a sort* is simply a given. It is the sort of naturalism that Quine famously articulates with Neurath’s metaphor of the gradual, piecemeal reconstruction of a boat already at sea:

The naturalistic philosopher begins his reasoning within the inherited world theory as a going concern. He tentatively believes all of it, but believes also that some unidentified portions are wrong. He tries to improve, clarify, and understand the system from within. He is the busy sailor adrift on Neurath’s boat. [Quine 1975: 72]

This is the sort of naturalism that takes philosophy to be “continuous” with science. If philosophy is the process of making explicit our commitments and integrating them into a coherent whole, then philosophy is of course continuous with science; science provides us with much of the raw materials, the body of commitments, which are the object of philosophical scrutiny. But philosophy is also continuous with everything else: so-called *properly basic* beliefs—the deliverances of sense-perception, memory and the

like—common sense, intuition, religious beliefs and practice, testimonial beliefs, and so on. This is not to say that all these things must have the same status for us; the philosophical process of sifting through commitments and weighing them against each other, may force one to revise and clarify her beliefs in various ways. If Neptune and Saturn are aligned and *yet* no promotion is forthcoming, then one's astrological beliefs may have to go. Nor is it to say that philosophers employ the same kinds of *methods* as scientists—drawing out commitments, incompatibilities, seeing if they can be integrated across different fields of thought is primarily an *a prioristic* affair, whereas scientific practice is inherently empirical. But the results of scientists are available to philosophers, and their methods themselves can be objects of philosophical enquiry. To ignore scientific practice and results when philosophising would be plain bizarre, and to take philosophy to undergird science, in the foundationalist sense, is an unwarranted and unmotivated extra step. Of course everyone, to a greater or lesser degree, considers the consequences, presuppositions and incompatibilities of their beliefs; this is not the sole province of the philosopher. Scientists may wonder what general epistemic principles may be latent in their information-gathering practices, people with religious commitments may think hard about what these commitments presuppose. Yet, when this practice is carried out *systematically*, it is what we call 'philosophy'.

1.2 The Credo

With all this firmly in mind, we shall take as our starting point David Lewis' 'credo', in which inveighs against nominalism as a rejection of mathematics:

Renouncing classes means rejecting mathematics. That will not do. Mathematics is an established, going concern. Philosophy is as shaky as can be. To reject mathematics on philosophical grounds would be absurd. If philosophers are sorely puzzled by the classes that constitute mathematical reality, that's our problem. We shouldn't expect mathematics to go away and make our lives easier. Even if we reject mathematics gently—explaining how it can be a most useful fiction, good without being true—we still reject it, and that's still absurd. [...] I laugh at how presumptuous it would be to reject mathematics for philosophical reasons. How would you like to go and tell the mathematicians that they must change their ways, and abjure countless errors, now that philosophy has discovered that there are no classes? Will you tell them, with a straight face, to follow philosophical argument wherever it leads? If they challenge your credentials, will you boast of philosophy's other great discoveries: that motion is impossible, that a being than which no greater can be conceived cannot be conceived not to exist, that it is unthinkable that anything exists outside the mind, ... and so on, and on ad nauseam? Not me! [Lewis 1991, reiterated in 1993]³

Lewis, disappointingly, fails to mention philosophy's other great discovery: concrete possible worlds, but we will forgive him this omission. It is instructive to consider Lewis' credo because I believe that

³Cf. Paseau [2005] for a discussion and disputation of Lewis' inductive pessimism argument against philosophical claims.

it contains a succinct account of why many platonists reject nominalism from the outset. As a result of the philosophical practice of looking inferentially upstream at what is presupposed by our various commitments, it is easy to think that philosophy has some kind of foundational role, in the sense that the practices of other disciplines await philosophical sanction in order to be validated. This is not so. Of course, it is always possible that by unearthing the presuppositions of some practice we find it to be in conflict with other ways of thinking, or with commitments that we are loath to give up or warranted in holding. In these cases philosophising can provide us with grounds for jettisoning commitments or practices we once endorsed. The conception of philosophy as foundational however goes beyond this; appending to this picture the notion that a special set of philosophical first principles have some kind of privileged role in our theorising, and will trump any other commitments seen to be in conflict with them. Perhaps Lewis takes nominalists to be in the grip of such a foundationalist conception of philosophy: nominalism is a philosophical first principle, and mathematics, both pure and applied, are in conflict with this principle and must be rejected. We have scientific evidence for the existence of mathematical objects, but this is trumped by philosophical evidence that they do not exist. This however is a serious misrepresentation of how contemporary nominalists actually reason. In fact, the landscape of the debate is better understood in the following way: both nominalists and platonists take mathematics to be an ‘established, going concern’, but platonists hold that our mathematical and scientific commitments rationally commit us also to the existence of mathematical objects. Nominalists deny this. The conflict then is not between mathematics and philosophy—a conflict that, Lewis assumes, mathematics is bound to win—but between two philosophical claims. Both the platonist and the nominalist make claims about what is rationally presupposed by a certain constellation of practices—the platonist that mathematical practices presuppose the existence of mathematical objects, and the nominalist that they do not—and claims of this sort are distinctively philosophical.

Not only does Lewis’ credo provide a succinct account of why many reject nominalism from the outset, it also provides a succinct account of what I take the task of the nominalist to be. Mathematics is a ‘going concern’, and it is the task of the nominalist to provide an account of this going concern that does not presuppose the existence of abstract objects. The goal of the philosopher is not to change mathematical practice; although perhaps there is no *a priori* reason to think that philosophical scrutiny could never, even *in principle*, warrant changing practices in mathematics or other respectable domains of inquiry. The goal of the philosopher is to account for mathematical practice *as it actually takes place*. At any rate, this will be my goal, and it is task enough to be getting on with.

1.3 Nominalism and Pragmatism

Over the last thirty years or so there has been a change in the character of the strategies espoused by philosophers actively defending nominalism, from what I will call *paraphrastic* nominalism—the attempt to avoid apparently problematic reference to and quantification over abstract objects by find-

ing ways to paraphrase away such usage—to non-paraphrastic nominalism. Charles Chihara's [1990] project was to replace the existential quantifier $\exists x$ with a *constructibility* quantifier Cx . One can replace talk of which mathematical objects exist, with talk of which open sentence tokens are constructible. Similarly, in Geoffrey Hellman's [1989] paraphrastic project, ordinary number-theoretic claims about what is *true* of the mathematical objects of Peano Arithmetic (PA) are taken to be elliptical for statements claiming (i) that there might have been a structure satisfying PA and (ii) that the claim *would* hold of that structure. Harry Field's [1980] paraphrastic programme has received by far the most attention in the philosophical literature. Field takes mathematical claims at face value, and holds them to be false. Instead of paraphrasing away pure mathematics, Field opts to motivate the plausibility of paraphrasing away *applied* mathematics—i.e. mathematicized theories of physics—to provide nominalistic counterpart theories which characterize the concrete domain in the same way as the mathematicized theories and have the same nominalistic consequences.

Latterly, nominalists who have been actively defending the doctrine have forsworn the paraphrastic route. In particular Jody Azzouni [2004] and Mary Leng [2010] have advocated forms of nominalism that do not involve any *reconstruction* of theories that refer to and quantify over mathematical objects. The lynchpin of Azzouni's approach is his claim that the existential quantifier $\exists x$ does not express existence. Hence, one can quantify over mathematical objects both in pure mathematics and in mathematicized scientific theories without incurring ontological commitment to mathematical objects. One can say true things about mathematical objects without those mathematical objects existing, as such the problems thought to accrue to the nominalist's commitment to the literal *falsity* of mathematical claims do not arise. Leng on the other hand accepts an existentially committing existential quantifier, but claims that one can make sense of the utility of reference to and quantification over mathematical objects without supposing that any such objects exist. Leng's goal is to produce a robustly *naturalistic* defence of nominalism: she argues that this claim can be defended making use of the epistemic standards inherent in scientific practice *itself*. (It is one of the strengths of this approach that a case for nominalism is spun from such parsimonious materials.) It has often been noted that non-paraphrastic approaches make the nominalist's task an easier one, but, as we will see, I think that the transition is a real philosophical, as opposed to merely *tactical*, advance on the part of the nominalist. Non-paraphrastic nominalism, so I say, cuts the issues at the joints better than its elder cousin; it reflects an improved understanding of the topic at hand.

Superficially, it would appear that the projects of Azzouni and Leng have more in common with each other than the old-school of Field, Chihara and Hellman, but this I think is a mistake. The nominalisms of Field, Chihara, Hellman and Azzouni are all akin in their insistence that one must, in the end, be in possession of scientific theories that are *true*, if one is, firstly to maintain that scientific theories really do furnish us with a means of accurately characterizing the world and, secondly, to account for the *utility* of scientific theorizing. Leng stands apart from rest of the ontologically parsimonious genus by having taken what I regard to be a *pragmatic* turn. Rather than assuming from the outset that *truth* must, at

bottom, be pivotal in accounting for the going concern of mathematics, one begins with mathematical practice and works backwards. If the truth of some claims that we are able to articulate is required to account for *this* then so be it, if not then so much the better for the nominalist. Whilst Leng herself does not explicitly characterize her project in these terms, a unifying theme running through these chapters, and the approach that I take it nominalists (and everyone else) ought to adopt, is that of *pragmatism*. This is not the *alethic* pragmatism of James [1907] in which a belief is true if and only if it is expedient to believe it, or indeed the more moderate alethic pragmatism of Peirce [1878] or Putnam [1981] in which a belief is true if and only if it would be accepted at an idealised limit point of inquiry. Rather, it is a kind of *methodological* pragmatism, or pragmatist order of explanation: we begin with how mathematics is used, what can be achieved through the use of mathematics, and from there seek to determine whether these facts of use require a distinctive domain of mathematical objects. *Doings* are given explanatory primacy; both in the sense that it is what we *do* with mathematical practices that needs to be accounted for, and in the sense that the roles of mathematics—for instance its role in representing concrete systems—are often best understood in terms of *practices* rather than in terms of the obtention of metaphysical facts. Lewis was correct: mathematics is a going concern. But this going concern is a cluster of practices, and these practices are what a philosophy of mathematics ought to account for.

1.4 Outlining the Project

I have said that I will provide a pragmatist defence of nominalism, but now it is time to lay out with more specificity the shape that this defence will take. In the second chapter I will characterize the target of nominalism: abstract objects. Abstract objects as I will understand them are objects that lack causal powers, that do not make a difference to the way other things are. I examine some challenges to the abstractness of mathematical objects, in particular an argument of Alan Baker's, and reject them.

In the third chapter I outline what I take to be the epistemological objection to abstract objects. The acausality of abstract objects results in severe epistemological problems for those who claim to have knowledge of their existence. This however is not because we ought to adopt a hefty causal theory of knowledge, but is a result of the platitude that knowledge excludes luck.

In the fourth chapter I address whether the existential quantifier should be taken to express existence.⁴ I examine three lines of thought that have been taken as evidence that the existential quantifier does not express existence, and will argue that none of the features of quantification they appeal to cannot be accommodated by an ontologically committing account of quantification. Accommodating these

⁴The debate has to do with the ring of a post-Fregean, analytic dispute, but its roots are in fact very old, stretching back at least to eighth-century Islamic theology. The Mu'tazilī's maintained that the category *thing* was divisible into the subcategories of *existent* and *nonexistent*. This was because contemporaneous Arabic grammarians held that 'thing' picks out anything that can be the subject of a predicate, and because of Qur'ānic verses in which God is described as creating the world by uttering "Be!" to a thing, which is henceforth *existent*. (Cf. Wisnovsky [2004] for discussion.) I however will be discussing some contemporary takes on the issue; although it is intriguing to see how similar were the considerations of medieval thinkers to contemporary ones.

features however involves accepting a notion of *nominalistic adequacy*, which in turn, I argue, can be put to use to circumvent an objection that nominalists cannot account for certain instances of quantification over mathematical objects which are used to facilitate inferences regarding concreta.

In the fifth chapter we turn to the indispensability argument and some of the issues affiliated with it. The relationship between truth and realism is considered. It is argued that nominalists can be realists in the sense that the nominalist can provide an account of how scientific theories can accurately characterize the world without being strictly true. The *applicability* of mathematics and the relationship of this topic to scientific representation, is discussed. I conclude that the application of mathematics provides no special problems for the nominalist. The explanatory role of mathematics is also considered. I develop an inferentialist account of understanding based on semantic inferentialist accounts of meaning, and define a minimal notion explanation in terms of it. I conclude also that the explanatory role of mathematics provides no special problems for the nominalist. Finally, Quine's method of discerning ontological commitment is criticized, and a new method of discerning ontological commitment is suggested.

In the final main chapter I consider the prospects of rolling out the nominalistic account of mathematics developed in the earlier chapters to accommodate other kinds of abstract objects. I use the test cases of *propositions* and *types*, and argue that the practices involving reference to and quantification over propositions and types can be explained nominalistically, in a manner cognate with mathematical practices. The prospects for a *unified* nominalism, I conclude, are good.

Chapter 2

Abstract Objects

Mathematical objects, if there are such things, do not cause anything to happen and, as such, do not make a difference. So stands the orthodoxy. Yet this orthodoxy faces a series of challenges. One lies in the conceptual gap between causation and difference-making. It may be granted that mathematical objects stand in no causal relations, but that there are other ways to make a difference. Such views may be bolstered by a wholesale rejection of causation as a now defunct artefact of folk metaphysics: fundamental physics shows us that there is no causal substratum, keeping observed regularities in place. Two such views will be considered here. The issue is of import to metaphysicians not only because it bears upon our understanding of what kind of things mathematical objects might be, but also because it bears upon whether they exist and how we have epistemic access to them. If mathematical objects make a difference—if it can be seen that things are the way they are *because* of mathematical objects—then we have a knock-down argument for their existence; nothing that doesn't exist can make a difference to what does. If, on the other hand, mathematical objects do not make a difference, the way is paved for a number of criticisms of platonism, not least that the fact seems incompatible with our gaining knowledge of their existence.

2.1 Mathematical Objects and Causation

I say that abstract objects do not exist, but what is it for an object to be abstract? It is often thought that there is some difficulty in analyzing the abstract-concrete distinction. If, for instance, one takes 'abstract' to mean 'non-physical', then one is burdened with explicating 'physical' in an acceptable way. Here I will simply stipulate what I mean by 'abstract' and 'concrete', although clearly any such stipulation, if it is to be useful, must bear a significant resemblance to the common parlance. Abstract objects, in my usage, are distinguished by their causal inefficacy. Gideon Rosen [2012] suggests a dilemma for the philosopher who opts to define abstract objects in this way. On the one hand, if one holds that an

object is abstract in virtue of being causally inefficacious, then one must count epiphenomenal qualia as abstract, as, although they are caused by brain states—or perhaps Boolean combinations of neuro-physical properties—they have no ‘downstream’ causal consequences of their own. Yet, epiphenomenal qualia are not ordinarily regarded as abstract. If, on the other hand, one avoids this consequence by strengthening the condition so that abstract objects are neither causes nor effects, then there are counterexamples; for impure sets of concrete objects, if such things exist, are effected by contingent facts about concreta. One can alter the cardinality of the set of origami cranes by creating an origami crane from a sheet of paper, or even bring about its annihilation by destroying all origami cranes, and bring it back into existence by folding a new crane. Rather than engage in terminological wrangling, I will simply define objects that are not causes or effects, such as pure sets, as *strongly* abstract, and objects that are not causes as *weakly* abstract. I am happy to say that epiphenomenal qualia are abstract in this sense.¹ At a first pass then, we will define abstractness in the following way:

[Weak Abstractness]: An object is *weakly abstract* iff_{df} it is not a cause of anything.

[Strong Abstractness]: An object is *strongly abstract* iff_{df} it is not a cause or effect of anything.

Abstract objects are objects that are either weakly abstract or strongly abstract. The categories *abstract* and *concrete* are exhaustive; an object is concrete iff_{df} it is not abstract.

I have defined abstractness in terms of causation, but it has been claimed that there is no such thing. Why reject causation? We might follow Russell—*circa* 1913, for Russell later abandoned the view—in rejecting causation on the grounds that fundamental physics gives us our best picture of reality, and that causation has no place in the picture supplied by fundamental physics. As Russell notes:

In the motions of mutually gravitating bodies, there is nothing that can be called a cause, and nothing that can be called an effect; there is merely a formula. Certain differential equations can be found, which hold at every instant for every particle of the system, and which, given the configuration and velocities at one instant, or the configurations at two instants, render the configuration at any other earlier or later instant theoretically calculable. That is to say, the configuration at any instant is a function of that instant and the configurations at two given instants. This statement holds throughout physics, and not only in the special case of gravitation. But there is nothing that could be properly called “cause” and nothing that could be properly called “effect” in such a system. [Russell 1913: 14]

This Russellian view may seem overly austere and prescriptive. Even if it is agreed that causation is not a fundamental constituent of reality, this may not in itself impel philosophers to abandon it wholesale. We may, for instance, note that although causation is absent from fundamental physics it abounds in the special sciences, which proceed largely by providing causal explanations for phenomena (and that those engaged in fundamental physics frequently help themselves also to a notion of causation when talking *about* fundamental physical theories, even if ‘causation’ does not appear *within* those theories).

¹Were epiphenomenal qualia to exist they would be importantly different from mathematical objects, even weakly abstract ones such as sets of concreta, insofar as they are *mind-dependent*. I do not, it should be noted, think plausible the claim that epiphenomenal qualia exist however.

To demand alterations to the practices of the special sciences on philosophical grounds would amount to a pointless trammelling of progress in those domains. Causal explanations are genuine explanations; it really is illuminating to learn that light stimulus to retinal ganglion cells causes a change in the rate of the action potentials they undergo. The same considerations would seem to apply to philosophy itself, for if the scientist can find profitable use of a notion of causation why not the epistemologist, or the philosopher of mathematics? If a full scientific understanding of the world necessarily makes appeal to causation then might not a full philosophical understanding of the world do the same? After all, not all of philosophy concerns itself with describing the fundamental nomic or counterfactual lineaments of the world. An eliminativist metaphysics of causation does not then, at least in any straightforward way, dictate the role of causation in philosophy more generally.

With this in mind we should be hesitant to cast the question of whether mathematical objects are causally active to the scrapheap of vacuous philosophical verbiage. Instead, the way to proceed is to determine firstly if there are useful construals of causation, and secondly whether it might reasonably be asked of mathematical objects stand in causal relations in the specified sense. One option is, with Ross, Ladyman and Spurrett [2007: ch.5], to take the identification of causes in the special sciences to amount to the identification of information-bearing real patterns. The leading idea here is that a pattern S is real iff (i) it is projectible—i.e., if there are measurable properties x_L and y_L of S , such that there is some physically possible way of effecting a (non-trivial) computation which, given x_L as input, yields y_L —and (ii) S has a model M which has less logical depth than a bit-map encoding of S , and where S is not projectible by a model of lower logical depth than M . (Cf. Ross, Ladyman and Spurrett [ibid.: ch.4]). We can then recast the question of whether mathematical objects are abstract in terms of whether mathematical objects play a role in information-bearing real patterns. It seems the answer is ‘no. Certainly scientists proceed by describing mathematical models in order to *represent* information-bearing patterns found in concrete systems, but they do not suppose that mathematical objects are *part* of those systems. Another option is to take the entire back light-cone of an event to be its cause.² An object partly causes an event if it is part of the entire back light-cone of an event. While this is not a particularly useful notion of causation for the special sciences, it is apposite for our purposes; for in identifying causes with entire back light cones, causation is reduced to something that *is* part of current physics—circumventing the Russellian objection—and the entire back light-cone of an event plausibly constitutes the broadest possible explanation, for any given event, of why it took place. Yet, it would be a category error to describe mathematical objects as being in a light-cone, so mathematical objects are acausal in this sense also.

Perhaps even if it is conceded that mathematical objects do not cause any events to take place, however we construe causation, one might worry that the focus has been too narrow. Mathematical objects may not be involved in causal relations, but this does not entail that they make no difference *tout court*. An object that makes a difference to the concrete world is not abstract in the sense I wish to capture. Another

²Maudlin [2007: ch.5] tables this idea.

aspect of our definition requires refinement. An object may not make a difference to anything because it is the only object that exists. A universe that contained a single mereological simple would not be a universe in which the mereological simple made a difference to anything else. Yet, the simple may not be abstract in the intuitive sense. In light of this, let us refine our definitions, in the first place so that they are not restricted to causal difference-making, and in the second place to avoid counterexamples involving concrete objects that are *contingently* unable to effect change on other objects:

[Weak Abstractness]: An object is *weakly abstract* iff_{df} it cannot make a difference to anything else.

[Strong Abstractness]: An object is *strongly abstract* iff_{df} nothing can make a difference to it, and it cannot make a difference to anything else.

There are at least two proposals that mathematical objects do make a non-causal difference, two proposals that mathematical objects are not abstract, which we will now take up.

2.2 Mathematical Objects and Conservativeness

One suggestion is that nonconservativeness tracks difference-making. A mathematical theory—*viz.* a theory that quantifies over mathematical objects—*S* is conservative over a nominalistic theory *N* iff for any nominalistic sentence *A*, if *A* is a consequence of *S+N* then it is a consequence of *N* alone. It may be thought that, if a body of sentences *S* is nonconservative over another body of sentences *N*—and both *S* and *N* are true—then the objects quantified over in *S* make a difference to the objects quantified over in *N*. The thought is this: *A*, not being a consequence of *N*, must be true *because* of *S*, or in virtue of *S*. But if this is so, then the truth of *S* makes a difference to the concrete domain, and since *S* is (partially) about mathematical objects, then mathematical objects make a difference to the concrete domain. The issue may appear otiose; has not Field, after all, proven that mathematics is conservative? The situation however is not so straightforward. Field's proof is for an applicable version of set theory he calls $ZFC_{V(N)}$. This is set theory with urelements, and with comprehension and replacement schemas, allowing us to form sets of kinds of concrete objects. Whether the conservativeness proof holds depends on the structure of the nominalistic theory *N*. If *N* contains an axiom scheme then applied set theory is only guaranteed to be conservative if the axiom scheme is stated as a list (of the form $\neg Q_0, \neg Q_1, \neg Q_2, \dots$). If the axiom scheme is stated as a rule however (for every formula *Q*, $\neg Q$ is an axiom), then applied set theory may be non-conservative over *N*.³

Most applied mathematics doesn't involve taking the union of a body of nominalistic claims and the form of applied set theory specified by Field, and making deductions. Most applied mathematics involves representing physical systems *as* mathematical structures. Theories such as this would simply be theories of pure mathematics but for the fact that its terms are associated with measurement procedures,

³Cf. Burgess and Rosen [1997: 194-6] for discussion.

and we take ourselves to be measuring aspects of concrete systems. When physical theories are couched in wholly mathematical language, the issue of conservativeness does not arise. What most philosophers have in mind however, when they discuss the conservativeness of mathematics is the addition of some mathematical theory, with bridge laws quantifying over both mathematical and concrete objects, to some accepted body of nominalistic claims. However, that mathematicized science is not normally carried out like this, does not diminish the philosophical import of the example. If nonconservativeness tracks difference-making and some mathematical theory is nonconservative over a nominalistic one, then mathematical objects make a difference.

What then of the claim that nonconservativeness tracks difference-making? We can see that something has gone amiss if we attend to specific examples. Take the nominalistic claim

(N*) There are planets

along with the mixed mathematical claim

(S*) There exists a bijection between the positive integers and the planets.

From S^*+N^* —but not from N^* alone—we can infer

(A*) There are infinitely many planets.

Here we can infer new nominalistic conclusions by forming the union of S^* and N^* , and some purely nominalistic facts are made true by S^* (on the assumption that S^* is true); yet it is clear that, in this case at least, the integers are not making a difference to the planets. The problem is that we cannot move from the *de dicto* claim that the truth of sentences about mathematical objects make a difference to the concrete domain, to the *de re* claim that mathematical objects *themselves* make a difference to the concrete domain. The reverse manoeuvre—from the *de re* claim that a class of objects make a difference to a domain D , to the *de dicto* claim that the truth of a sentence about that class of objects makes a difference to D —is legitimate: if an object makes a difference then a full and accurate description of the world must depict those differences. As we have seen however, the nonconservativeness of applied mathematics is not sufficient to show that mathematical objects make a difference.

2.3 Mathematical Objects and Laws

Alan Baker [2003] has argued that whether mathematical objects make a difference to the concrete domain is dependent on whether mathematics is indispensable for the formulation of scientific laws: specifically that *if* mathematics is dispensable in the formulation of physical laws then mathematical objects make no difference to the concrete domain, but if mathematics is indispensable in the formulation of physical laws then it is unclear whether mathematical objects make a difference to the concrete domain, ‘perhaps irredeemably so’. Baker’s argument is original and worth separate examination, and proceeds by considering how we are to assess the counterfactual:

(No-Difference) If there were no mathematical objects then (according to platonism) this would make no difference to the concrete, physical world. [Baker 2003, p225]

How we assess this counterfactual, says Baker, depends on whether mathematics is indispensable to the formulation of physical laws. But why?

There are two ways in which we can assess counterfactuals. Stalnaker's [1968] method takes $\phi \Box \rightarrow \psi$ to hold at a world w iff *if* ϕ is true at some world x —and for any world y such that ϕ is true, and x is at least as close to w as y —then ψ is true at x . Or, more demotically, $\phi \Box \rightarrow \psi$ is true iff in the closest possible world in which ϕ is true, ψ is also true. Lewis' [1973] method takes $\phi \Box \rightarrow \psi$ to hold iff either ϕ is true at no worlds, or there is a world x in which ϕ is true, and for all worlds y if y is at least as similar to the actual world as x then $\phi \rightarrow \psi$ is true at y .⁴ Baker opts for the Stalnaker method, so the question then becomes this: *Given a world in which mathematical objects exist, is it the case that in the closest possible world in which mathematical objects do not exist the concrete domain remains unchanged?*

It may seem obvious that the answer to this is affirmative; after all, if we think of the world as consisting of an abstract and a concrete domain, then the closest world without an abstract domain is the one in which the concrete domain is identical to our own (Baker calls this world PW). Baker suggests an interesting way in which we might cast doubt on this, by appealing to how we understand the closeness of worlds. Although there is no workable objective measure of similarity, Lewis suggests (and Baker endorses), as a metric for measuring the closeness of possible worlds when evaluating counterfactuals, the following criteria:

1. It is of first importance to avoid big, widespread, diverse violations of law.
2. It is of second importance to maximize the spatiotemporal region throughout which perfect match of particular fact prevails.
3. It is of third importance to avoid even small, localized, simple violations of law.
4. It is of little or no importance to secure approximate similarity of particular facts, even in matters that concern us greatly. [Lewis 1979, p442]

Given Lewis' criteria for measuring the closeness of possible worlds, we arrive at the following premiss:

⁴Philosophers engage in a good deal of hand-wringing rumination regarding our ability to assess counterfactuals with impossible antecedents, or the coherence of doing so. This I think is misguided; we have no trouble assessing counterfactuals such as *If my parents were different I would have had a different upbringing*. Nor is there anything conceptually incoherent in imagining mathematical objects ceasing to exist, or never existing in the first place—a point ably defended by Field [1993]. I might reasonably be thought that counterfactuals such as 'if the number 7 had not existed then...' are of dubious intelligibility. This however is grounded in the structuralist intuition that it does not make sense to think of *individual* numbers existing or failing to exist, and is not a mark against the intelligibility of counterfactuals regarding the existence or non-existence of numbers *in toto*. What *does* merit concern is the use of metaphysically possible worlds to assess counterfactuals, when there are legitimate uses of counterfactuals involving the impossible. But here the proper way of proceeding is to amend the apparatus to fit the task at hand, not to stymie the task at hand to fit the apparatus; and the amendment involves moving from talk of metaphysically possible worlds to talk of epistemically or conceptually possible worlds.

- (1) For all worlds w_i, w_j , if the laws of nature in a world w_i are different than in a world w_j , then w_i is distant from w_j .

Whether the use of mathematics in formulating physical laws is dispensable, Baker argues, affects whether PW involves widespread violations of laws that obtain in the actual world. If quantification over mathematical objects is eliminable from our best scientific theories, then we can describe the physical laws that obtain purely ‘nominalistically’. In this case the laws that are true in our world are true in PW. On the other hand, if quantification over mathematical objects is ineliminable from our best scientific theories, then the laws that are true in the actual world will not be true in PW, quantifying as they do over mathematical objects that do not exist there. This provides Baker’s second key premiss:

- (2) For all worlds w_i , if w_i is a world in which mathematical objects do not exist, the laws of nature at w_i are different to the actual world.⁵

Given (1) and (2), Baker argues, the ability to dispense with mathematical language would imply that (No-Difference) holds, while if we are unable to dispense with mathematical language, the story is a little more complicated. Certainly PW would not be a close possible world, as it has different physical laws from the actual world; but whether there are any *closer* worlds in which mathematical objects do not exist looks like something we don’t have the resources to assess. For instance, assuming that Newtonian physics is nominalizable, it may be that a Newtonian world is closer to the actual than PW, given that Newtonian mechanics is at least approximately true in the actual world. Baker’s point however, is that it isn’t clear that there is any principled basis on which to adjudicate in cases such as these.

Strange that our ability to *articulate* the regularities that hold in our world without resorting to quantification over mathematical objects should have this kind of bearing on which counterfactuals are true, but this result falls out of a particular way of interpreting the notion of a *law of nature*. Laws are sentences that define (abstract) models⁶ and are trivially true of the models they characterise. Laws are ascribed to the world via theoretical hypotheses which are constituted by claims that concrete systems are relevantly similar (e.g. isomorphic, homomorphic) to the models that represent them. But this is not the only way we think about laws. *That* no signals travel faster than light is a law of nature, but it is not a sentence. Moreover it features in explanations of, for instance, why our images of Alpha Centauri can only show us what was taking place there more than four years ago. We claim that scientists seek to discover laws of nature or perhaps that facts about the mental supervene on physical facts and laws, and that God did not need to supplement these with psychological laws. When we speak this way, we refer not to sentences that characterise models, but to the nomic necessities or regularities that we take to obtain in the actual world—the regularities that we use abstract models to represent.⁷ There are then

⁵The assumption here that mathematical objects exist at the actual world is not question-begging in this context, since the goal is to assess the counterfactual (No-Difference).

⁶Although there are other ways of defining models, for instance concrete scale models.

⁷It may be contended that laws in this sense are propositions, given that they are denoted by *that*-clauses. Contrary to the rather pervasive assumption that *that*-clauses always denote propositions, there are tests employed by linguists which show that this is not generally the case. Referential noun phrases admit unconstrained substitution in extensional contexts, so that if *that*-clauses denoted propositions then it would be possible to replace *that*-clauses with nominal constructions of the form *the proposition that S*. Yet, we cannot. For instance, ‘Peter wishes that we be sober-minded’ cannot be rendered ‘Peter wishes the proposition

two notions of law which we might appeal to: a linguistic notion, and a metaphysical notion. The two notions are related: the abstract models which laws (linguistic) characterise are models *of* concrete systems in which the laws (metaphysical) obtain.⁸ In discussions of laws, philosophers often move freely between these two notions, and for the most part the practice is harmless enough, but in this case it has a bearing on the soundness of Baker's argument.

If one follows Baker in assessing counterfactuals by appealing to laws in the linguistic sense, then one runs into problems. Firstly, it does not follow that, if mathematics is dispensable from the formulation of our best scientific theories, that PW does not involve widespread violations of law. Rather, we find ourselves with two sets of true laws, one nominalistic, the other mathematicized, and the latter is still violated. Contrast this with the metaphysical conception of laws. On the metaphysical conception we have not two different sets of laws, but two *different means of characterizing* the *same* laws of nature; one that quantifies over mathematical objects and one that does not. Secondly, 'laws of nature', in the sense used (and *required*) by Baker, only exist if the denizens of the world have articulated them; a world with the same nomic necessities/regularities as our own, but populated by ultra-articulate aliens who have described those regularities without mathematical vocabulary, has different laws of nature to our world. Thirdly, if there is no way of nominalistically formulating our best physical laws, then PW has *no laws of nature*; which is to say that it has the same laws as other lawless worlds, including those where events simply take place at random. But it is absurd in this context to claim that worlds with the same observable regularities as our own have the same laws as worlds in which events take place at random, and are in that sense not importantly dissimilar. The lesson here is that, the appropriate sense of 'law' to employ when assessing counterfactuals is the metaphysical one.

Let us now assess the two key premises of Baker's argument:

- (1) For all worlds w_i, w_j , if the laws of nature in a world w_i are different than in a world w_j , then w_i is distant from w_j .
- (2) For all worlds w_i , if w_i is a world in which mathematical objects do not exist, the laws of nature at w_i are different to the actual world.

Given (1) and (2) Baker infers that PW is distant from the actual world and hence not clearly the closest world in which mathematical objects do not exist. But the argument is faced with a trilemma:

Suppose that we interpret 'laws', as Baker does, in the linguistic sense. In this case (2) is true, so long as no nominalization of the laws of nature has been carried out. (1) however is false; since two worlds could exhibit precisely the same regularities, governed by the same nomic necessities, match each other exactly in particular fact, and yet have different laws in the linguistic sense. As I have argued, the linguistic notion of a law of nature is ill-suited to the assessment of counterfactuals. Suppose that instead

that we be sober-minded'. Closer to home, while 'That no signals travel faster than light, prevents us from communication with extra-terrestrials' makes sense, 'The proposition that no signals travel faster than light, prevents us from communication with extra-terrestrials' does not. Cf. Moltmann [2003] for discussion.

⁸There is a heavy-duty platonist view described by Field [1989, ch6] which takes the practice of mathematicized science *not* to involve representing concrete systems as models, but rather describing mixed mathematical-physical objects *directly*, without the intermediary of a model. This view I will discuss, and reject, in chapter 5.

we interpret laws in the metaphysical sense. In this case (1) is true, but (2) is false, since mathematical and nominalistic theories aim to characterize the same nomic necessities or regularities that hold in a world. Finally, one can secure the truth of both premisses by understanding ‘laws’ metaphysically in (1) and linguistically in (2), in which case both (1) and (2) are true, but the argument involves an equivocation. In any case, the argument does not go through. Moreover, once one thinks of laws as the nomic necessities or regularities that obtain in a world—as I have claimed we ought, when assessing counterfactuals—it is clear that the closest world in which mathematical objects do not exist is the world in which the same concrete objects exist and are governed by the same nomic necessities or regularities as the actual world; i.e. PW. As such, mathematical objects do not make a difference.

Chapter 3

Epistemic Luck and Necessary Truths

3.1 Prologue

The most common objection to platonism is epistemological, although what the epistemological objection *is* commands less concurrence. This is my attempt to make out that objection. While I do not think that abstract objects exist, it is worth emphasizing that the objection I raise here is *de jure* rather than *de facto*. I will not present arguments that mathematical objects do not exist but rather that belief in them does not and cannot constitute knowledge. That one cannot possess knowledge of mathematical objects is however reason enough to withhold belief in them. The epistemological problem with platonism is, I say, serious, yet there have been various attempts to nip it in the bud. Lewis, for instance, attempts to defuse the objection by casting doubt on the reasonableness of demanding an explanation of how evidence for abstract objects could possibly be garnered. He states:

I think it is true that causal acquaintance is required for some sorts of knowledge but not for others. However, the department of knowledge that requires causal acquaintance is not demarcated by its concrete subject matter. It is demarcated instead by its contingency. Here, the relevance is plain. If I know by seeing, for instance, my visual experience depends on the scene before my eyes; if the scene had been different, within limits, my experience and my subsequent belief would have been correspondingly different. [...] But nothing can depend counterfactually on non-contingent matters. For instance, nothing can depend counterfactually on what mathematical objects there are... Nothing sensible can be said about how our opinions would be different if there were no number seventeen [Lewis 1986: 111]

Field describes Lewis as trying to show that the lack of an explanation of the reliability of our mathematical beliefs should not be upsetting, and devotes a number of pages attempting to counter the argument.

[Field 1989: 233-9]. However, taken by itself Lewis' argument amounts to nothing more than a re-statement of the original epistemic problem; it is only by a feat of rhetoric that he manages to make it appear otherwise. The nominalist will concur that one cannot know of mathematical objects through counterfactual contingencies and that it makes no sense to try to, but will still ask how one *can* know about them.¹ Burgess makes a different kind of gambit:

The epistemological argument, according to which belief in abstract objects, even if conceded to be implicit in scientific and commonsense thought, and even if perhaps true – for the aim of going epistemological is precisely to avoid direct confrontation over the question of the truth of anti-nominalist existence claims – cannot constitute knowledge, surely is not intended as a Gettierological observation about the gap between justified true belief and what may properly be called knowledge. It follows that it must be an issue about justification; and here to the naturalized anti-nominalist the nominalist appears simply to be substituting some extra-, supra-, praeter-scientific philosophical standard of justification for the ordinary standards of justification employed by science and common sense. [Burgess 2008: 5]

The tactic here is to make the nominalist appear committed to a kind of Cartesian foundationalism; in particular, one that is forced to reject ordinary scientific standards of justification. It is not clear why one might think that that epistemological objection to platonism ‘surely is not intended as a Gettierological observation about the gap between justified true belief and what may properly be called knowledge’, and we might never know, because Burgess isn't telling; the pronouncement is left unargued. At any rate, the argument I present here *is* about the gap between justified true belief and knowledge. To criticize belief in abstract objects, or for that matter almost anything, on the grounds that it is not justified is problematic, because justification is agent-relative. Almost anyone can be justified in believing almost anything, so long as the right sorts of conditions are in place. Six-day creationism may be justified for a person who believes it on the testimony of otherwise reliable interlocutors, and has no exposure to evidence that might undermine the doctrine. Moreover, it is surely uncharitable to claim that no-one is justified in taking abstract objects to exist. Many platonist philosophers have carefully and diligently studied the arguments for and against the position, have come to the conclusion that there are intractable problems with nominalism, and have adopted platonism accordingly.

My hope is to break some of the deadlocks that have to an extent immobilized the debate. Stalemate has set in largely for the reason that nominalists have failed to provide an *argument* that knowledge of mathematical and other abstract objects cannot be had, and have been content to stress the vague, if somewhat forceful, intuition that there is *something* amiss about the idea that we could come to gain knowledge of a domain of objects to which we could not, even in principle, interact with. In order to redress this state of affairs—and attempt, probably in vain—to break the stalemate, I present a deductive argument, proceeding from platitudes about knowledge, that knowledge of abstract objects

¹To give Lewis his due, he does not appear to intend the statement to be an argument *for* platonistic mathematical knowledge. In fact, he goes on to say that we ought to believe in abstract objects because of the role they play in theory which is ‘systematic and devoid of arbitrariness’ [Lewis 1986: 111]. Of course, to pursue this point would take us deep into the territory of indispensability and no-miracle arguments, and this is a topic we shall discuss later.

is an impossibility. The platitude at hand is that knowledge excludes luck. It has been argued that the modal ‘safety’ condition on knowledge acts to rule out epistemic luck, and so should be treated as a necessary condition for knowledge. However, at least one class of beliefs—belief in propositions which are necessarily true—cannot be accommodated by the safety condition, as necessary truths trivially meet such modal requirements by being true in every possible world. I consider the received method in accommodating such cases, and provide counterexamples. I then argue that recent work in epistemic possibility provides resources for extending the notion of epistemic luck to cover cases of necessary truths. I go on to show how such an approach handles other counterexamples to safety and sheds new light on the epistemological problem with mathematical objects.

3.2 The Luck Platitude

Anti-luck epistemology has witnessed a flurry of activity in recent years. The platitude that knowledge in some sense excludes luck, has long been central to epistemological theorizing, and perhaps more explicitly so since Gettier cases became part of the common currency of epistemology, yet has only recently received the attention it deserves.² This has enkindled the project of anti-luck epistemology, which takes epistemic luck as being central to our understanding of knowledge. This chapter is both a criticism and an endorsement of this particular brand of epistemology: a criticism insofar as I believe anti-luck epistemology is ill equipped to deal with an important class of knowledge; an endorsement insofar as I think it can be made, albeit with some slight modification, to do so. In what follows I will give a sketch of the anti-luck requirement on knowledge as well as the safety condition on knowledge which is thought to manifest this requirement. I will then argue that, as the case of luckily true belief in necessary propositions illustrates, the safety condition is incapable in principle of capturing the epistemic luck requirement, offer a replacement safety condition, and show how it can be used to shed light, not only on epistemic luck, but on the epistemological problem with abstract objects.

We have said that knowledge excludes luck, but some refinements have to be made. Knowledge does not exclude luck *simpliciter*. It may, for instance, be lucky that the proposition in question is true. Perhaps you are a lottery winner. This does not mean that you cannot have knowledge that you are a lottery winner; you may have excellent grounds for your belief. So, *content epistemic luck*—when it is lucky that the proposition is true—is benign. So too is *capacity epistemic luck*—when it is lucky that the agent is capable of knowledge. Perhaps you have made a remarkable recovery from a head injury that ought to have rendered you incapable of thought. You are lucky to be capable of knowledge, but capable of knowledge nonetheless. Another benign form of luck is *evidential epistemic luck*—when it is lucky that you acquire the evidence on which you base your belief. Perhaps you are a detective who discovers, quite by chance, a suspect’s DNA at a crime scene. The element of luck here does not prevent the detective from achieving knowledge.

²Cf. Pritchard 2005 for a book-length treatment.

The malign form of epistemic luck is known as *veritic epistemic luck*—when it is lucky that the agents belief is true. Pritchard offers the following modal condition for veritically lucky true belief:

S's true belief is lucky *iff* there is a wide class of near-by possible worlds in which S continues to believe the target proposition, and the relevant initial conditions for the formation of that belief are the same as in the actual world, and yet the belief is false. [Pritchard 2007, p280]

3.2.1 Safety

If knowledge excludes veritic luck then we require some kind of condition that eliminates veritic luck. Here, safety³ seems to fit the bill:

S's true belief p is safe *iff* in nearly all (if not all) nearby possible worlds w in which S believes p , p is true.

This needs some refinement. Borrowing an example of Sosa's [cf. Sosa 2007, p26], suppose I am hit hard and experience significant pain. On the basis of being hit I believe, and know, that I am in pain. Suppose also that I am a hypochondriac and would have believed myself to be in pain even if I had suffered only a very slight blow. My belief is not safe, as I would likely have held it even if it were false. As such we require the following amendment:

S's belief p is safe *iff* in nearly all (if not all) nearby possible worlds w in which S forms a belief p on the same basis as in the actual world, p is true.

The basis then for my belief is that of being struck hard. Those worlds in which I suffer only a very light blow need not be taken into consideration. If knowledge excludes veritic luck and the amended safety condition eliminates veritic luck, then the amended safety condition—or something of its ilk—is a necessary condition for knowledge.⁴

3.2.2 The Problem

Our modal condition is designed to eliminate epistemic luck. However it seems incapable of doing so when the propositions in question are necessarily true. A quick glance at possible world semantics exposes the problem. Consider the following clause:

(Nec) $v(\Box p, w) = T$ *iff* for every world w' in W , such that Rww' , $v(p, w') = T$

where $v(p, w)$ is the truth value of p at world w , W is the set of possible worlds, and R is an accessibility relation on worlds. This states that $\Box p$ is true at a world w when p is true at all (accessible) possible

³Proposed by Sosa [1999] who parsed the condition as 'a belief by S that p [is safe] *iff*: S would believe that p only if it were so that p ' [p142].

⁴Note that even *Robust Virtue Epistemologists* such as Sosa, Greco and Zagzebski, who wish to parse knowledge purely in terms of the cognitive abilities of agents, must accept some anti-luck condition on knowledge, or otherwise reject the platitude that knowledge excludes luck. In this case, whatever their favoured virtue-theoretic conditions for knowledge, these must at least implicitly contain an anti-luck condition; *viz.* these conditions must *entail* an anti-luck condition.

worlds. From (Nec) it follows that:

(Nec-Safe) If $\Box p$, then in all nearby possible worlds w in which S forms a belief p on the same basis as in the actual world, p is true.

In other words, if p is necessarily true then S 's belief p is (trivially) safe. But surely there are cases of luckily true belief in necessary propositions. Consider Gullible Joe who comes to believe (correctly) that his parents are cousins on the basis of an insult directed at him. Let us say, furthermore, that the person directing the insult knows nothing about Joe. If, with Kripke, we think that people have their parents essentially, then Gullible Joe's belief is necessarily true, and, as such, trivially safe. Yet, Gullible Joe does not have knowledge that his parents are cousins. Moreover, it seems that Joe does not have knowledge as a result of his belief being only luckily true. So, there are cases where the safety condition does not achieve its end of eliminating epistemic luck. Joe's belief is lucky, but the world modally over-cooperates and his belief is safe. If we are to deal with such cases the safety condition will have to be extended or ameliorated in some way.

At this juncture one option is simply to restrict our account to contingent propositions, but even this is not entirely sufficient. Suppose that Gullible Joe forms the belief that objects contract when in motion, subsequent to his having taken a powerful hallucinogenic. There are no nearby possible worlds in which Lorentz Contraction does not take place—as such worlds are nomologically impossible—thus Joe's belief is safe. So, if we are to restrict the account then we must go further and restrict it to 'fully contingent' propositions: propositions which are nomologically, as well as metaphysically or logically possible. This option is not appealing; it would be desirable to have an account of epistemic luck which was applicable to all propositions. Joe's belief here seems luckily true; a universally applicable account of epistemic luck could shed light on what is common to all cases of luckily true belief.

I've claimed that we need to modify or extend our account of epistemic luck in some way in order to accommodate luck pertaining to necessary truths. However, there is already an account of how this might be achieved:

all we need to do is to talk of the doxastic result of the target belief-forming process, whatever that might be, and not focus solely on belief in the target proposition. For example, if one forms one's belief that $2 + 2 = 4$ by tossing a coin, then while there are no near-by possible worlds where that belief is false, there is a wide class of near-by possible worlds where that belief-forming process brings about a doxastic result which is false (e.g., a possible world in which one in this way forms the belief that $2 + 2 = 5$). The focus on fully contingent propositions is thus simply a way of simplifying the account; it does not represent an admission that the account only applies to a restricted class of propositions. [Pritchard 2008, p3]⁵

⁵This approach is taken from Mišćević [2007] where he suggests a requirement of *Agent Stability* for necessarily true propositions: 'If an agent knows a priori a (necessary) proposition p , then, in most nearby possible worlds in which she forms her belief about p in a slightly different way or with slightly changed cognitive apparatus as in the actual world, that agent will also come to believe that p .' [ibid. p.62] In fact, this kind of response is pre-empted by Sainsbury, who states 'It is easily possible for me to be wrong in believing that p (even if it is true that p) iff at some world close to the actual world the actual episode of forming the belief that p (or the counterpart for this episode) is one in which a false belief is formed. In some cases, this is because the same

It may be thought that a deficiency of this approach is that it provides a bifurcate account of epistemic luck: S's belief that p is non-lucky iff it is safe, unless there are no nearby possible worlds in which $\neg p$, in which case... As we have considered, there seems to be something in common to all cases of luckily true belief, and only a unified account of epistemic luck could capture this. However, the criterion could be formulated so as to handle both kinds of cases. Weatherson contrasts *Content-safety* where 'B is safe iff p is true in all similar worlds' and *Belief-safety* where 'B is safe iff B is true in all similar worlds' [Weatherson 2004, p378]. Incorporating the notion of a basis to avoid the kind of counterexamples discussed above we might endorse:

(Safety*) S's belief is safe iff in all nearby possible worlds in which S forms a belief on the same basis as in the actual world S forms a true belief.

Safety* provides a unified account of luck in both cases where the proposition believed is contingent and where it is necessary. A more pressing problem though—both for the disjunctive account, and for Safety*—is that they rely on the 'doxastic result of the target belief-forming process' being somehow unstable; but this need not be the case. There is nothing to preclude examples where (i) what is believed is necessary, (ii) the proposition believed is fixed across close possible worlds, and (iii) the belief is luckily true. Consider Gullible Joe, who forms beliefs about logical truths by shaking a lucky 8-ball which, instead of predicting the future, only states logical truths.

It is worth reflecting here on the project we are engaged in. Given that safety is normally presented as a necessary rather than sufficient condition for knowledge, our objection to it may seem wide of the mark. We may hope to explain the failure to gain knowledge by appealing to some other principle. In this case neither reliability nor 'sensitivity' will do the trick, for Joe's belief-forming processes reliably lead him to these necessarily true beliefs, and sensitivity is trivially met in cases of necessarily true propositions for similar reasons to safety (a subjunctive conditional with a necessarily false antecedent is trivially true). Nor will justification help in all cases. Note that we need not suppose that Joe is epistemically blameworthy or irrational in any way; he may be the victim of a widespread and sophisticated collusion. But even if there were some condition which acted as a band-aid for cases such as these, to invoke it would miss the point. Joe is *lucky* that his beliefs are true; an account of epistemic luck should try to capture this. That safety does not however is not, in itself, a reason to reject it. It is plausible that no modal condition can ever be sufficient for knowledge. If an absence of epistemic luck is indeed all that is required to elevate true belief into knowledge, then the flip side of this would be that no modal condition can ever be necessary for epistemic luck. But this is no counsel of despair. Even if we cannot *fully* analyse epistemic luck with a modal condition, this should not prevent us from trying to capture more of the contours of epistemic luck with an improved modal condition. To think the project thus useless is to commit a kind of fallacy of perfection. A Final Analysis of knowledge can instead be thought of as the limit point towards which our efforts tend. My aim here then, is to offer an improved account of belief is, but is false at the close alternative world. In other cases, it is because a different, and false, belief is formed at the close world. [Sainsbury 1997 p.913]

epistemic luck, if not a perfect one. As we will see, improving our modal condition allows us to put epistemic luck to work in ways that were previously precluded.

3.3 Epistemic Possibility

We have, up until this point, been dealing with counterfactuals, but there is another sense of possibility—epistemic possibility—which may be of more use to us. Epistemic possibility deals with the way things might be for all we know (this is to be contrasted with counterfactual or metaphysical possibility, which deals with ways things might have been). So, for all I know, it is possible that some form of string theory is true, or that Goldbach’s conjecture is correct (or incorrect). For all I know for certain, it is possible that I was adopted, that water is not H₂O, or that intuitionistic logic is sound. Underlying, and conceptually antecedent to, this notion, which Chalmers [2010] calls *strict epistemic possibility*, is the notion of *deep epistemic possibility*—ways things might be, prior to what anyone knows.

How we choose to carve up epistemic space depends on our purposes. If, for the purposes of simplicity and tractability, we wish to model idealized agents then we may consider all propositions which can, in principle, be ruled out *a priori* as epistemically impossible. Perhaps all scenarios which involve logical contradictions would be one such class of epistemically impossible worlds, or more generally, what can be known *a priori* is epistemically necessary. When modeling non-idealized agents we will have to weaken these constraints somewhat; for instance, the truth of naïve set theory may be epistemically possible for an agent who is unaware of, and not in a position to discover, Russell’s paradox. Rather than talking about what is epistemically possible *per se* we can consider what is epistemically possible for an agent S at time *t*. One way to accommodate this would be to allow that *a priori* knowable propositions are not epistemically necessary for agents who have not yet come to know them *a priori*. For example, that ‘ $p \rightarrow q$ ’ contraposes with ‘ $\neg q \rightarrow \neg p$ ’ would not be epistemically necessary for someone who is logically competent, but who was yet to give the matter thought.

One characteristic feature of epistemic possibility, is that what is epistemically possible may not be metaphysically possible. It’s epistemically possible, for all I know for certain, that water is not H₂O, yet, it is metaphysically impossible that water is not H₂O. So too, it’s epistemically possible that my parents are not who I think they are, but if they *are* who I think they are, then it’s metaphysically impossible that they could be anyone else. To avoid confusion, we ought not to talk of possible worlds when we deal with epistemic possibility as, on the usual understanding, there are no possible worlds in which, say, Hesperus is not Phosphorus, yet it is epistemically possible for many agents that Hesperus is not Phosphorus. As such, we talk of scenarios which *verify* or *falsify* sentences or propositions: a scenario *w* verifies a sentence *s* when the sentence is true in that scenario, and falsifies a sentence when the sentence is false in that scenario. How should we think of scenarios?

Clearly we cannot straightforwardly identify them with possible worlds for the reason just mentioned:

some scenarios are metaphysically impossible. However, this problem can be evaded, still using the apparatus of metaphysically possible worlds, by distinguishing between *verification* and *satisfaction*, and by making scenarios ‘centered’ worlds, where a centered world is an ordered triple $\langle w, \phi_1, \phi_2 \rangle$ of a (metaphysically possible) world, an individual and a time. In this way we augment our third-person description of a world with indexical information regarding the world’s centre. If an individual x is at the centre of a world w then we can make use of a predicate ϕ which is only true of x at w . Our ordered triple then, will be a complete objective description of a world, along with sentences of the form ‘I am ϕ_1 ’ and ‘Now is ϕ_2 ’. Even though no metaphysically possible world *satisfies* ‘Water is XYZ’—which is to say, there are no scenarios where it is *true* that water is XYZ—we can now say (roughly) that a scenario *verifies* ‘Water is XYZ’ if it is a centered world where ‘water’ picks out the watery stuff with which the individual at the centre of the world is familiar, which, at this world, is XYZ. Issues still remain, a particularly salient one being ‘strong necessities’; necessary truths that are verified by all centered worlds. If God exists necessarily, for instance, then there will be no centered worlds where God does not exist. Yet, it seems that it is deeply epistemically possible that God does not exist. Some too equate nomological necessity with metaphysical necessity. If such a view is correct then there will be no centered worlds which exemplify different physical laws from our own; yet it seems that it is deeply epistemically possible that, for instance, Lorentz contraction or time dilation do not take place.

A better strategy is to bypass such thorny issues and construct epistemic scenarios from the ground up, independent of possible worlds. For even if we concur with those, like Chalmers [2002], who deny that there are any strong necessities, epistemic possibility is orthogonal to such metaphysical issues, and our account of epistemic possibility should reflect this autonomy. Instead then of identifying scenarios with triples of the form $\langle w, \phi_1, \phi_2 \rangle$, as we considered above, we can think of scenarios as sentence types of an ideal language \mathcal{L} . This language will permit infinite sentences, in order to describe some kinds of infinite scenarios. It must also use only ‘epistemically invariant’ expressions: expressions whose epistemic import will not change from context to context, or utterance to utterance. It is often thought that names and natural kind terms are not epistemically invariant (apart from perhaps in the mouth of God), nor are context-sensitive terms. A *fully specified* scenario can be identified with a sentence d of \mathcal{L} which is epistemically possible and for which there is no sentence s such that $d \& s$ and $d \& \neg s$ are both epistemically possible.⁶

What should we say about the metaphysics of scenarios? In a discussion of *impossible* worlds, Graham Priest [1997] suggests a parity thesis: whatever theory we have of possible worlds can be brought to bear on impossible worlds, as there are no good reasons to suppose that there should be an ontological difference between possible and impossible worlds. If one is a modal realist with regards to the former then so too should they be a modal realist with regards to the latter; if possible worlds are states of affairs

⁶Of course the same could be achieved without scenarios by extending ordinary possible world semantics to include (metaphysically) impossible worlds. However, the only impossible worlds that will be of relevance when assessing safety will be those that are epistemically possible, so talk of epistemically possible scenarios is more apposite.

(à la Plantinga 1974), ways the world could have been (à la Stalnaker 1976)⁷, non-existent objects (à la Priest 2005), or mere fictions (à la Rosen 1990), then so are impossible worlds. A similar parity thesis suggests itself with regards to scenarios; however, as we will see, it may be difficult to maintain a metaphysics of scenarios other than fictionalism, or the noneism of Priest.

Having given this brief sketch of epistemic possibility,⁸ its relevance should be clear. That the metaphysically impossible is often epistemically possible means that epistemic possibility may be better placed as a tool for analyzing epistemic luck, as it may help us provide an account of veritic luck regarding belief in metaphysically necessary propositions. We might suggest the following reformulation of the safety principle:

Epistemic Safety:

S's belief p is safe iff in nearly all (if not all) nearby⁹ scenarios w in which S forms a belief p on the same basis as the actual world, p is true in w .

This new formulation has a number of advantages. For one thing, it allows us to see how veritic luck can infect belief in necessary propositions: if a proposition p happens to be metaphysically necessary, S's belief p isn't automatically trivially safe.

3.3.1 Refinements

However, more work needs to be done; refinements to the epistemic safety principle are required. It seems that some epistemically possible scenarios—scenarios which are nomologically impossible or which involve massive changes in particular fact—are still too far away for lucky beliefs to be unsafe. Consider an agent who forms a luckily true belief that water is H_2O . Are there really close possible scenarios in which water is not H_2O ? Scenarios in which the liquid which occupies 71% of the earth's surface is different to the actual world would be rather distant, or, at any rate, too far away to be troubled by a safety requirement. As such, the account needs to be finessed.

One option is to adopt a different modal requirement—either in place of epistemic safety or as a supplement to it—for instance an epistemic counterpart to tracking conditions such as Dretske's 'conclusive reasons' [Dretske 1971] or the sensitivity condition suggested by Nozick [1981]:

Epistemic Sensitivity:

S's belief p is sensitive iff in the closest scenario or scenarios, in which S forms their belief in the same way as in the actual world, and in which $\neg p$, S does not believe p .

Using epistemic sensitivity as our modal requirement is my favoured option. This deals neatly with the above example, and, although sensitivity is subject to a number of criticisms, it isn't clear to me that they cannot be overcome. For instance, the sensitivity condition is often rejected as it violates the closure

⁷Soames [2007] adopts this metaphysics with regards to scenarios.

⁸Cf. e.g. Chalmers [2006 and, especially, 2010] for a more worked-out account of epistemic possibility.

⁹The distance of these scenarios is measured in the usual way, suggested by Lewis [1979, p472], with respect to the actual world.

of knowledge under known entailment. One knows that one has hands and knows that this entails that one is not a brain in a vat, yet the belief that one is not a brain in a vat is not sensitive and hence not knowledge. Whilst Nozick's original sensitivity condition indeed has this consequence—a consequence that he embraced—epistemic sensitivity avoids this by stipulating that S forms her belief in the same way as in the actual world. It is often also thought that a sensitivity condition is too onerous a condition for knowledge. This thought is motivated by scenarios such as the following:

On my way to the elevator I release a trash bag down the chute from my high rise condo. Presumably I know my bag will soon be in the basement. But what if, having been released, it still (incredibly) were not to arrive there? That presumably would be because it had been snagged somehow in the chute on the way down (an incredibly rare occurrence), or some such happenstance. But none such could affect my predictive belief as I release it, so I would still predict that the bag would soon arrive in the basement. My belief seems not to be sensitive, therefore, but constitutes knowledge anyhow, and can correctly be said to do so. [Sosa 1999: 145-6]

The counterfactual is quite difficult to assess here. Perhaps the closest possible world in which the bag is snagged is one in which garbage chutes are quite generally less reliable than they are in the actual world, in which case one may well not form the belief that it has reached the bottom. Setting this complication aside, whilst these thought experiments lie in penumbral areas of our intuitions, I am happy to deny that, in the ordinary case, a person *knows* that their refuse has reached the bottom of the garbage chute. One's intuitions about knowledge are easily cross-wired with other epistemic intuitions, when contemplating scenarios such as these. For instance, it is clear that the person here is highly *justified* in their belief, and even *knows* a number of propositions in the vicinity; for instance that it is *highly probable* that their garbage has reached the bottom of the chute. Moreover, a sensitivity condition can be motivated in the following way: Suppose that S holds a belief *p* on a particular basis *b*. S is then informed that even if *p* were false, *b* would still hold. It seems clear that S has a *defeater* for her belief *p*; that she should, in light of this new information, give up her belief *p*, or, at least, her claim to know *p*. But if sensitivity is not a necessary condition for knowledge then it is not clear why this should be so. A comprehensive discussion of these issues would take us too far afield; nevertheless, I note my demurral from the mainstream. However, I think that safety theorists can also circumvent these issues by eschewing ordering scenarios relative to the actual world—which can be thought of as a fully-specified scenario—in favour of ordering scenarios relative to a partially specified scenario *w'* which verifies only the set $\{\psi : \psi \in \Sigma(S, @)\}$ of propositions ψ specifying S's basis for belief at the actual world @. In cases where an agent's belief that water is H₂O is based on the role of a dice, for instance, *w'* would be 'agnostic' with regards to the chemical composition of water; hence there would be many nearby scenarios in which water is not H₂O. In this way, we could reformulate the epistemic safety condition as:

Revised Epistemic Safety:

S's belief *p* is safe iff in nearly all (if not all) scenarios *w*—nearby to the partial scenario

w' , which verifies $\{\psi : \psi \in \Sigma(S, @)\}$ —in which S forms a belief p on the same basis as the actual world, p is true in w .

It is important to realize that this is not an internalist reconstrual of safety; there is no reason to equate S 's basis for belief with her evidence at $@$, where 'evidence' is thought of internalistically. Handling necessary truths by retreating to internalism would significantly diminish the interest of the approach, but, worse still, doing so would risk scepticism, as brains in vats being fed experiences could possess equivalent bases to agents in normal epistemic environments. Nor can we remedy this by following Williamson [2000, ch.9] in equating evidence with knowledge, on pain of circularity. Rather, we should continue to think of bases in precisely the way that is implicit in traditional basis-relative accounts of safety. Bases in these accounts are partly constituted by internal facts about the agent's beliefs and experiences, but also consist of those events which bring about the agent's belief; including external events (such as being struck hard). Exactly how bases should be delineated is an important question, and getting it right will be crucial for the viability of any safety condition on knowledge, but this is not something we will address here: firstly because it is a topic demanding a fairly substantive discussion on its own; but also because it is an issue for *everyone* who adopts a basis-relative modal condition for knowledge, not just the advocate of epistemic safety (although epistemic safety serves to draw attention to the fact that it is an issue).¹⁰ The intuitive notion of a basis will, in any case, be adequate for our purposes.

The epistemologist's aim is to model epistemic luck, in much the same way that a physicist may model fluid mechanics or a biologist the population growth of a species. Our methodology too is akin to that of the scientist; we construct a formal model used to represent the object of investigation, examine the logical consequences of that model and verify that it accords with the data at hand. In doing so we hope to unearth the deep structure that lies behind the phenomena we are investigating. But there are reasons to think that the traditional model of possible worlds could never capture important structural features of epistemic luck. For one thing, by dealing only with possible worlds, traditional safety builds in an assumption that only possibly true propositions are epistemically relevant to agents. But we often find ourselves in situations where propositions that are metaphysically impossible are on the table; for instance if we are trying to work out who our biological parents are, from a group of (epistemically) possible candidates. For another thing, just as the world can be such that (luckily) it makes our beliefs true—as in the Gettier cases—and just as the world can be such that (luckily) it makes our beliefs reliable—as in Plantinga's brain lesion case¹¹—so too, the world can be such that (luckily) it makes our beliefs safe. In each instance, features of the world, irrelevant to any cognitive activity on the part of

¹⁰Here is just one matter that would need to be resolved. Part of the basis for belief-formation are other beliefs held by the agent (a fact attested to by the numerous examples in the psychology literature of how prior beliefs can act to condition the beliefs that we go on to form). An obvious requirement then would be that the beliefs that we take into consideration when assessing safety (in both its traditional and epistemic incarnations) are true or approximately true. This condition would rule out cases of safe-but-lucky belief in which the agent forms a belief p on the basis of false beliefs. This itself raises questions: the prior beliefs of the agent may be true, but luckily so; in which case we will have another instance of safe-but-lucky beliefs. Something then must be said about the aetiology of the true beliefs that are to be included in S 's basis. To reiterate: these are questions for *all* theorists who adopt some modal condition for knowledge.

¹¹Discussed in section 3.2.

the agent, collude to make their belief true. By shifting the emphasis away from the world, to the world *insofar as it is presented to and impinges upon the agent*, the epistemic safety condition better reflects the structure of epistemic luck.

3.4 Putting Epistemic Safety to Work

3.4.1 Irrelevance

One happy result of epistemic safety is that it deals nicely with a class of counterexamples to the analysis of epistemic luck in terms of safety, suggested by Jennifer Lackey [2008]. Lackey puts forward the following example:

‘consider [a person, Penelope,] winning through a lucky guess a game show that presents contestants with multiple choice options. Now imagine that there is a feature, ϕ , of the final winning answer that is entirely disconnected from its correctness but is such that its presence will invariably lead to Penelope to choose that answer. Suppose further that the current producer of the show, Gustaf, has a similar obsession with ϕ , so that he ensures that the final winning answer of the day will possess this feature. Perhaps ϕ is being presented in the color purple, so that when in doubt Penelope will invariably choose the answer displayed in purple and Gustaf will always present the final winning answer in purple.’ [Lackey 2008: 263]

Despite being a paradigmatically lucky event, Penelope’s guess is safe, as there are no close possible worlds in which Gustaf does not present the winning answer in purple and no close possible worlds in which Penelope picks an answer which is not presented in purple. The fact of Gustaf’s purple fixation ‘*just happen[s] to fortuitously combine*’ with Penelope’s similar obsession, to render the event safe. Lackey suggests a recipe for constructing such counterexamples:

first choose a paradigmatic instance of luck, such as winning a game show through a purely lucky guess, emerging unharmed from an otherwise fatal accident through no special assistance, etc. Second, construct a case in which, though both central aspects of the event are counterfactually robust, there is no deliberate or otherwise relevant connection between them. Third, if there are any residual doubts that such an event [is safe] add further features to guarantee counterfactual robustness across nearby possible worlds. [ibid.]

Here, epistemic safety can be put to use. Recall that Epistemic Safety orders scenarios relative to the agent’s basis for belief— $\{\psi : \psi \in \Sigma(S, @)\}$ —at the actual world. Gustaf’s obsession with purple is not part of Penelope’s basis for the belief that the correct answer will be presented in purple- this fact is outwith Penelope’s purview. Hence, there are many scenarios close to $\{\psi : \psi \in \Sigma(S, @)\}$ in which Gustaf does not present the winning answer in purple and, as such, many nearby scenarios in which Penelope answers incorrectly. Epistemic safety gives the correct result that picking the correct answer is lucky for Penelope, *given her basis for belief*.

We have seen that traditional safety is too weak in one respect; it cannot capture luck in cases of necessary truths or where irrelevant factors act to fortuitously fix the belief and the fact believed across close possible worlds. But there are also cases where safety is too strong as a necessary condition for knowledge. Juan Comesaña suggests the following example:

HALLOWEEN PARTY: There is a Halloween party at Andy's house, and I am invited. Andy's house is very difficult to find, so he hires Judy to stand at a crossroads and direct people towards the house (Judy's job is to tell people that the party is at the house down the left road). Unbeknownst to me, Andy doesn't want Michael to go to the party, so he also tells Judy that if she sees Michael she should tell him the same thing she tells everybody else (that the party is at the house down the left road), but she should immediately phone Andy so that the party can be moved to Adam's house, which is down the right road. I seriously consider disguising myself as Michael, but at the last moment I don't. When I get to the crossroads, I ask Judy where the party is, and she tells me that it is down the left road. [Comesaña 2005: 397]

Despite knowing where the party is, my belief isn't safe, so the objection goes, because there is a close possible world in which I disguise myself as Michael and, hence, a close possible world where I form a belief that the party is at Andy's house on the same basis as the actual world—Judy's testimony that the party is at the house down the left road—but where my belief is false. Here the possibility that I dress as Michael and am subsequently misled is not antithetical to my *actually* gaining knowledge, yet, because it takes place in a close possible world, traditional safety infelicitously factors it in. Cases such as these¹² work by envisaging a situation in which the agent has the kind of connection to the facts which ensures knowledge, and then loading that scenario with epistemically irrelevant features which are designed to ensure that there are close possible worlds in which the proposition believed by the agent is not true. Epistemic safety handles cases of this sort because in assessing the closeness of possible worlds it only factors in those events which constitute the agent's basis for belief, and ignores the epistemically irrelevant features which make the belief unsafe. In this case epistemic safety doesn't factor in Andy and Judy's collusion, as this is no part of my basis for belief. As such, the scenarios close to $\{\psi : \psi \in \Sigma(S, @)\}$ where I continue to believe that p on the same basis as the actual world, will not include situations where the party has been moved to Adam's house.

3.4.2 Mathematical Objects

It is often contended that there is something epistemically dubious about mathematical objects. Given their mind-independence, coupled with their acausal nature, there is a concern that, even if such entities exist, it seems impossible that we might gain knowledge of them. Despite this, it has not always been clear exactly how to frame this supposed epistemological problem. The canonical articulation of the

¹²Ram Neta and Guy Rohrbaugh [2004] suggest a counterexample to safety with the same structure as Comesaña's.

epistemological problem for platonism is of course Benacerraf's famous paper *Mathematical Truth*¹³, which has prompted a torrent of discussion since its publication. Benacerraf claimed that our account of mathematical knowledge 'must fit into an over-all account of knowledge in a way that makes it intelligible how we have the mathematical knowledge that we have' [Benacerraf 1973: 409]. It is well known both that the over-all account of knowledge Benacerraf had in mind was a causal theory of knowledge¹⁴ of some description and that causal theories of knowledge have since become unpopular;¹⁵ these are issues that will not be rehearsed here. What is of interest is that, absent a causal theory of knowledge, it is difficult to articulate what the problem might be; so much so that some recent work still assumes that, if there is to be an epistemic problem with abstract objects, it must somehow be grounded in a causal requirement for knowledge (cf., for example, Potter [2007] and Wetzel [2009]). Nonetheless, there is a felt sense that *something* is epistemically worrying about mathematical objects; a sense testified to by the large number of attempts by those who think we can have knowledge of mathematical objects, to provide some kind of explanation of how this is possible.

Field reformulates the challenge in terms of providing an account of how our mathematical beliefs are reliable:

The mathematical realist believes that his or her own states of mathematical belief, and those of most members of the mathematical community, are to a large extent disquotationally true. This means that those belief states are highly correlated with mathematical facts: more precisely (and without talk of truth or facts), that for most mathematical sentences that you substitute for 'p', the following holds:

If mathematicians accept 'p' then p.

[...] the fact that [this schema] hold[s] for the most part is surely a fact that requires explanation: we need an explanation of how it can have come about that mathematicians' belief states and utterances so well reflect the mathematical facts. [Field 1989: 230]

It is wrong to presume that Field's challenge presupposes a reliabilist account of knowledge. In fact, Field's challenge does not seem to presuppose *any* theory of knowledge; Field is careful not to mention knowledge, justification, or other such epistemic notions. Rather, the platonist's inability to explain the reliability of our mathematical beliefs is intended to be thought of as a problem in itself. Two issues arise: the first is that this approach is dialectically limited insofar as it is not an *argument against* knowledge of mathematical objects. Anti-nominalists will find nothing here that might impel them to concede that we cannot have knowledge of mathematical objects. The second stems from there being good reasons to think that reliability is not a sufficient condition for knowledge. Plantinga [1993] describes a case involving a patient whose brain lesion causes him to believe that he has a brain lesion.

¹³Benacerraf 1973

¹⁴First suggested by Alvin Goldman [1967].

¹⁵But by no means extinct. Colin Cheyne [1998] espouses a causal theory of existential knowledge using examples from empirical science to show that such a precedent has already been set. The argument is intended as a rebuttal to Bob Hale [1994], where Hale argues that any causal theory that would block knowledge of platonic objects would require that every known fact be causally connected to the knowers belief in that fact; which would in turn block knowledge of other facts. Cheyne avoids this criticism by limiting the causal criterion to existential knowledge. These views are more fully explicated in Cheyne [2001].

We are to suppose that the patient has no evidence for his belief whatsoever; even that he has evidence against having a lesion, but that the lesion prevents him from assimilating this information appropriately. The belief-forming process is reliable, but the resultant belief will, to use Plantinga's preferred term, lack *warrant*.

This leaves open the possibility that—as in Plantinga's brain lesion example—we could possess an account of the reliability of our mathematical beliefs which does nothing to solve the epistemological problem. Thus, Field's formulation of the challenge could be met, even when we have no account of how knowledge of mathematical objects is possible. Moreover, a mathematical parallel to the brain lesion case actually exists. Balaguer [1998] has shown that by simply increasing the content of the 'platonic realm', so as to include every consistent mathematical object, reliability can be achieved trivially. In rough outline: so long as every consistent mathematical object happens to exist, and so long as our mathematical beliefs are consistent, then they will always be true, and thus, reliable. Although this is presented as a genuine solution to the platonist's epistemological problem, I take it that "solving" the epistemological problems of a controversial ontology by increasing that ontology to its limit *ought* to be seen as a philosophical sleight of hand. Someone who gets the mathematical facts right merely on this basis no more has knowledge than Plantinga's brain lesion victim.

Modal requirements on knowledge have previously been thought inapplicable to the case of mathematical objects, as it is often contended that, if mathematical objects exist then they do so necessarily.¹⁶ However, the applicability of our revised safety requirement is not at the mercy of the modal status of mathematical objects. Recall the requirement:

S's belief *p* is safe iff in nearly all (if not all) scenarios *w*—nearby to the partial scenario *w'*, which verifies $\{\psi : \psi \in \Sigma(S, @)\}$ —in which S forms a belief *p* on the same basis as the actual world, *p* is true in *w*.

We now have a way of formulating the epistemological problem with mathematical objects. If we accept that:

1. Epistemic safety/sensitivity holds for a belief that *p* iff that belief is non-lucky;
2. Lucky beliefs are not knowledge; and
3. Beliefs about mathematical objects are epistemically unsafe;

then we are committed to beliefs about mathematical objects not being knowledge.

How confident should we be in 1–3? 2 seems to be on very strong ground: it is widely regarded to be a platitude that knowledge excludes (veritic) luck. The second question is whether our beliefs about mathematical objects are epistemically safe or sensitive. It is clear that they cannot be epistemically

¹⁶Not everyone agrees. Field [1993] argues compellingly for the contingency of mathematical objects, and this is also a natural, though not obligatory, view for Quinians (who view mathematical claims as a posteriori knowable theoretical claims) to adopt. This view has however remained heterodox.

sensitive. Suppose *S* believes some proposition *p* regarding mathematical objects. Given their acausal nature, mathematical objects can have no bearing on *S*'s basis for belief, so that in the closest scenarios in which *S* forms a belief on the same basis as the actual world and in which *p*, *S* would still believe *p*.¹⁷ A similar story applies to epistemic safety. *S*'s belief is epistemically safe when her basis for belief places the right kind of constraints on the way the world might turn out to be. But as we have noted, abstract objects can have no bearing on a person's basis for belief—nor can that basis for belief have any bearing on abstract objects—so that an agent's basis for belief can place no constraints on how things might turn out to be with abstract objects. As such, when *S* forms a belief *p*, regarding abstract objects, there will always be a large number of scenarios, nearby to the partial scenario which verifies *S*'s basis for belief, in which *S* forms the belief *p* on the same basis as in the actual world but in which *p* is false. Beliefs regarding mathematical objects are not epistemically safe.¹⁸ Importantly, although it's the acausality of mathematical objects that results in knowledge of their existence being impossible, this isn't because of some causal requirement on knowledge. Rather, it falls out of the platitude that knowledge excludes luck, once we spell out what this platitude amounts to.

With regards to 1,¹⁹ I have been making the case that epistemic safety or sensitivity excludes epistemic luck. Whether we accept this will depend on whether these conditions prove to be a faithful model of epistemic luck. Traditional safety offered a creditable account of luck, and acts to shed light on a number of issues in epistemology where luck plays a central role. Epistemic safety and sensitivity can also play this role, whilst extending it to cover cases of belief in necessary propositions, and dealing with counterexamples to the traditional rendering of these conditions. Given the fruitfulness of the approach, combined with its accordance with our intuitions regarding luck, it seems plausible that epistemic safety or sensitivity does play the role of excluding luck. Of course, it is always (epistemically) possible that a counterexample will be found. If this is indeed the case then the safety requirement will either have to be modified accordingly. Whether our beliefs about mathematical objects are deemed lucky will then depend on how the new requirement is formulated, or on the details of what is put in its place. However, even if this is the case, it is difficult to see how any such replacement would allow for knowledge of mathematical or other abstract objects. For there must be *some* modal condition on knowledge, if our account is to capture the incompatibility of knowledge with epistemic luck; and the acausality of abstract objects entails that there are *no* modal conditions on knowledge—one these modal conditions are properly understood to involve epistemic modalities—that belief in those objects could meet. More can always be said, but there is reason to think that epistemic safety, sensitivity, or some development thereof, is not only well placed to capture the anti-luck platitude with regards to knowledge, but also

¹⁷Some philosophers object to accessing conditionals with (potentially) metaphysically impossible antecedents on the grounds that we are being asked to countenance something which is unintelligible. However there seems to be nothing unintelligible in statements such as *If I had different parents then I would have been raised differently*, *If intuitionistic logic is correct then the law of excluded middle does not hold* or for that matter *If mathematical objects did not exist then...*

¹⁸We are now in a position to see why we must adopt a fictionalist or noneist stance to scenarios; a more ontologically robust view would be subject to precisely the same epistemological argument we have advanced against mathematical platonism.

¹⁹Note that we need not be committed to the biconditional 1 for the conclusion to go through. The weaker premiss *A belief that p is non-lucky only if it is epistemically safe/sensitive* would suffice.

to shed light on wider philosophical issues. Moreover, it seems clear that traditional counterfactual approaches, are, by their very nature, incapable of doing either.

Chapter 4

Existential Quantification: The Third Way

“What is there?”, said Quine, can be answered in one word—‘Everything’ and everyone will accept the answer as true’ [Quine 1948]. Or perhaps they won’t. Quine’s theory of quantification has recently been subject to an assault on two fronts, one stemming from debates over nominalism in the philosophy of mathematics, and the other from metaphysics, where the substantiveness of a number of ontological debates has been called into question. I will provide a sketch of the considerations that have inclined philosophers to doubt the standard approach to ontological commitment. I will then lay out some considerations in favour of the Quinian view of quantification. I will argue that these apparently competing considerations can be reconciled by shifting away from the traditional focus on the semantics of quantification and attending (also) to its pragmatics, arguing that in doing so we can delineate a new account of ontological commitment which capitalizes on the insights of current theories whilst circumventing their problems. Having developed a framework for understanding the relationship between ontological commitment and the existential quantifier, which draws on a notion of *adequacy*, we turn our attention to logical consequence and the ‘speed-up’ phenomenon, showing how this framework can be put to work in other areas of the philosophy of logic.

4.1 The Quinian Orthodoxy

Quine’s theory of ontological commitment has become firmly established as the orthodoxy, as such a résumé of his views are in order. Quine’s criterion finds its canonical statement in *On What There Is*:

To be assumed as an entity is ... to be reckoned as the value of a variable. [...] The variables of quantification, ‘something’, ‘nothing’, ‘everything’, range over our whole on-

tology, whatever it may be; and we are convicted of a particular ontological presupposition if, and only if, the alleged presuppositum has to be reckoned among the entities over which our variables range in order to render one of our affirmations true. [Quine 1961: 13]

So, for example, if the sentence ‘there are D-branes and they have mass’ is true, then D-branes exist. Despite its huge influence on the philosophical community, Quine’s paper offers no argument that the existential quantifier expresses ontological commitment [cf. Priest 2008], but this *can* be found elsewhere:

What there are, according to a given theory in standard form, are all and only the objects that the variables of quantification are meant in that theory to take as values. This is scarcely contestable, since “ (x) ” and “ $(\exists x)$ ” are explained by the words “each object x is such that” and “there is an object x such that”. Some other languages may have no clear equivalent of our existential phrase “there is”, nor of our quantifiers; but surely there is no putting the two assunder. [Quine 1986: 89]

Quine argues that our natural language locution ‘there is’ expresses ontological commitment and that the existential quantifier takes its meaning from the former. As such, the existential quantifier itself inherits this feature.

4.2 Doubting the Quinian Orthodoxy

4.2.1 Metametaphysics (I)

A kind of crisis of confidence looms over some prominent debates in metaphysics. We are confronted with a leporine arrangement of simples, but is there a further object, a rabbit, composed of these simples? Mereological ‘nihilists’ such as Cian Dorr [2005] think not: there are only particles arranged rabbit-wise. Universalists such as David Lewis [1986] or James Van Cleve [2008] think there is, and that for any objects there is another which has those objects as its parts. Yet others, such as Ned Markosian [2008], believe that composition sometimes takes place, perhaps only when a living creature is involved [van Inwagen 1990]. But for some there is a lingering suspicion that the whole debate is empty in some way; a verbal spat which has somehow donned the vestments of substantive debate.¹ Most metaphysicists think that such debates *are* substantial, and this has found a defence in Fine [2000, 2009], Sider [2001, 2009] and van Inwagen [1998, 2002, 2009]. This realism is characterized by the view that there are determinate answers to such ontological questions, and that these answers are substantial and to be

¹Kant’s exhortation merits repetition:

To know what questions may reasonably be asked is already a great and necessary proof of sagacity and insight. For if a question is absurd in itself and calls for an answer where none is required, it not only brings shame on the propounder of the question, but may betray an incautious listener into absurd answers, thus presenting, as the ancients said, the ludicrous spectacle of one man milking a he-goat and the other holding a sieve underneath. [Kant, *Critique of Pure Reason*, A57/B82-3]

unearthed by serious and sustained metaphysical work. Other authors retain the first requirement but jettison the second: there are objective and determinate answers to these questions, but these answers are not ‘metaphysically robust’—in some sense, usually left obscure—and are obvious or can be immediately inferred from obvious truths. Hale and Wright [2001, 2005, 2009], Hirsch [1993, 2009] and Thomasson [2007, 2009] have all defended views of this sort. Finally, some, such as Balaguer [1998], Chalmers (tentatively) [2009], Putnam [1987, 2004], Sidelle [2002] and Yablo [1998, 2009], reject both requirements and endorse an ontological anti-realism according to which there are no determinate answers to (at least some) ontological disputes.

But these issues seem to pertain to the nature of reality, not of language; whence then the debate over the existential quantifier? The claim that there are objects that are composed of other objects can be stated ‘ $\exists x \exists y \exists z (x = \text{sum}(y, z))$ ’, or the claim that there are rabbits ‘ $\exists x \text{rabbit}(x)$ ’. If we accept bivalence and the claim that the existential quantifier expresses existence, then we are committed to there being a determinate answer to the question of whether there exist rabbits or objects composed of other objects more generally. If one holds that the debate over composite objects is not substantive, but doesn’t wish to reject bivalence then the most viable recourse may be to treat the existential quantifier as non-ontologically committing.

4.2.2 Metametaphysics (II)

A second source of doubt emerging from metametaphysical concerns, articulated by Hofweber [2005, 2007], stems from the observation that there is a seeming mismatch between the substantiveness of ontological questions, and the trivial way in which we often seem to answer them. Consider for instance:

There are three divine persons.

∴ The number of divine persons is three.

∴ There is some thing that is the number of divine persons.

∴ There are numbers.

∴ Numbers exist.

Arguments of this sort appear valid due (i) to a commitment to Quine’s criterion, and (ii) to our tendency to treat certain nominalistic sentences as equivalent to seemingly more ontologically loaded counterparts, in this case:

The number of divine persons is three iff there are three divine persons.

which is an instance of a more general principle:

(Num) The number of F s = n iff there are n F s.

There is a strong sense that (Num) is correct, indeed analytically so; yet the left hand side of the biconditional appears committed to the existence of numbers whereas the right hand side does not. One way of reconciling the equivalence of the left and right sides of (Num) is to hold that the existence of mathematical objects is in some sense ‘insubstantial’. The implicit reasoning is *modus tollens*: If a metaphysical question is substantial then it cannot be answered trivially. The dispute over the existence of mathematical objects can be answered trivially, therefore it is not substantial. If the debate is insubstantial then we require some deflationary way of understanding it. Perhaps when we existentially quantify over numbers we do not express ‘robust’ existence (they are somehow second-class entities), or perhaps there is no fact of the matter whether numbers exist. Either response involves rejecting the existential quantifier as univocally expressing ontological commitment.

In fact matters are not nearly as clear as is often supposed. The common (Fregean) practice of understanding numerals as referring to abstract numbers in sentences such as ‘The number of BNP councillors is 28’ can lead to problems. For combined with the sentence ‘The number of BNP councillors is (used predicatively) appalling’ we can infer ‘28 is appalling’. Moltmann [forthcoming] argues that we should understand *number-of* terms as referring to number tropes possessed by pluralities of objects. The sorts of properties that can be legitimately attributed to number tropes are the sorts of properties that can be legitimately attributed to the plurality of entities in which the tropes reside. In the case at hand, the number of BNP councillors is the sort of thing that can be appalling because BNP councillors are the sort of thing that can be appalling. This view hews to the linguistic data better than the Fregean one, and, significantly, does not treat the left hand side of the biconditional as committed to an ontology to which the right hand side is not. Whilst it is plausible then that, contrary to the received view, that taking n in (Num) to range over abstract numbers does not correctly characterise our natural language use of *number-of* terms, we will assume presently that the Fregean reading of (Num) has some independent plausibility as a principle in its own right.

4.2.3 Linguistic Data

In the philosophy of mathematics, doubts about the existential quantifier’s ontological robustness come from two sources. The first is mathematical nominalism. Nominalists believe that abstract mathematical objects such as numbers, sets, functions and *ante rem* structures do not exist. Yet, our everyday discourse, not to mention our scientific claims, is replete with quantification over abstract objects. For instance, if the sentence ‘1729 is the smallest number which can be written as the sum of two cubes in two different ways’ is true then, given Quine’s criterion, numbers exist.

One option is to become an error theorist regarding such discourse. This view has some significant proponents—Field [1980, 1989], Balaguer (in nominalist mode) [1998] and Leng [2010]—but many find it unattractive. ‘Our mathematical assertions (when they go well) are surely true’, they say, ‘and we ought not to reject them for the sake of abstruse ontological scruples’. However, the claim that

our natural language locution ‘there is’ is ontologically committing has been challenged, primarily by Jody Azzouni [1997, 2004, 2006, 2007 2009, 2010] and Graham Priest [2005, 2008]. Their challenge emerges from observations of how existential quantification is used in ordinary language. Consider, for instance, the following paragraph:

(*) There are many ways to skin a cat. Thus, although Seymour had not a clue about the whereabouts of the leprechaun that he had read about when he was young, he realized that others—in more supernatural lines of work—might. It was his obsession with this particular leprechaun that had caused him to fail to complete his thesis, and that incompleted thesis, in turn, was why there were so many jobs he had never gotten (although he was entitled to them), nor why he had never built the house he had been planning to build for so many years. So Seymour had many motives for finding the whereabouts of that leprechaun. [Azzouni, 2007: 213]

The suggestion is that someone may utter this sentence without believing in or intending to convey the existence of *leprechauns*, *lines of work*, *theses*, *jobs*, *unbuilt houses*, *motives* and the like. Such uses of quantification demand explanation, a potential one being that the ‘there is’ idiom and the existential quantifier do not express existence. Priest draws out the tension between existential quantification and the expression of existence more explicitly still by noting the use of existential quantification in order to then express non-existence:

I thought of something I would like to buy you for Christmas, but I couldn’t get it because it doesn’t exist (e.g., a perpetual motion machine). [Priest 2008: 43]

Priest suggests calling the existential quantifier the ‘particular’ quantifier to avoid biasing the issue in favour of the Quinian orthodoxy, and to ossify this typographically introduces a new symbol $\exists x$ meaning ‘some x ’, whose domain ranges over non-existent as well as existent objects and, as such, doesn’t express existence. An \mathfrak{A} quantifier is also introduced ‘to keep \exists company’—even though it can be understood in precisely the same way as \forall —as is an existence predicate E .² At this juncture the Quinian will likely respond that the sentences proffered by Azzouni and Priest simply haven’t been properly regimented, and that an appropriate schematization into first-order logic will reveal their true ontological commitments. Quine himself recognized that our natural language locutions could, ontologically speaking, lead us astray:

In a loose way we often can speak of ontological presuppositions at the level of ordinary language, but this makes sense just in so far as we have in mind some likeliest, most obvious way of schematizing the discourse in question along quantificational lines. It is here that the ‘there is’ of ordinary English lends its services as a fallible guide—an all too fallible one if we pursue it purely as philologists, unmindful of the readiest routes of logical schematization.’ [1980: 107].

Thus, according to the Quinian, what *really* reveals to us our ontological commitments is a suitable regimentation of our English locutions into first-order logic. Here, two points present themselves. Firstly,

²We may think it cannot be the case that \mathfrak{A} is equivalent to \forall , as $\forall x$ is equivalent to $\neg\exists x\neg$, but for Priest the domain of quantification includes non-existent as well as existent objects.

this assumes that ontologically apposite regimentations are always to be had, but we are offered no reason to suppose that this must be the case. Secondly, this is in tension with the view that the existential quantifier inherits its meaning (and subsequently its expression of ontological commitment) from its natural language counterpart, because it appears to entail that the ‘there is’ idiom of English is not (at least always) meant to express existence; in which case the major premiss of Quine’s argument looks unmotivated. Thus the Quinian faces a dilemma: either she accept the English usage of ‘there is’ at face value and retain the claim that ‘ \exists ’ takes its meaning from it, whilst losing its ontological commitment, or she does not accept the English usage of ‘there is’ at face value, but loses the motivation to think that ‘ \exists ’ is ontologically committing.³⁴

4.3 All Roads Lead to Quine

We have briefly examined the sources of pressure on the Quinian orthodoxy, but what of the considerations in its favour? Quine’s own argument was found to be problematic, but perhaps there are other reasons to think that the existential quantifier expresses existence. In fact, I think that there are such reasons. At least two stem from constraints on an existential quantifier that is to play a useful role in formal languages.

4.3.1 Trivializing Truth-Conditions

There is much of interest in Priest’s treatment of non-existence as it pertains to intentionality [cf. Priest 2005], and we might hope that it would provide the materials for a semantics for the existential quantifier which could handle the examples of quantification in the vernacular examined above. However, it cannot do so, at least as it stands. Priest’s view is a version of Richard Routley’s [1980] *noneism*.⁵ We have already noted that Priest does not take the particular quantifier to express existence, but this raises questions over how we are to assess the truth of existentially quantified sentences; clearly there is nothing in a domain of extant objects which will satisfy sentences that quantify over that which doesn’t exist. How, for instance, would we assess the truth of ‘There are fictional mice that talk’?⁶ In Priest’s model, the domain D for every world w is the set of *all* objects, existent or non-existent. A formula $Pt_1\dots t_n$ is true at w iff $\langle \delta(t_1), \dots, \delta(t_n) \rangle \in \delta^+(P, w)$, where if c is a constant $\delta(c) \in D$, and $\delta^+(P, w)$ is the extension of P at w . A formula $\exists xA$ is true at w iff for some $d \in D$ $A(x/d)$ is true at w . Being tacit Quinians, philosophers tend to use variable-domain modal logics, capturing, as they do, the thought that what there is at one world may not exist at another. But, as Priest observes:

³We will return to this second response later.

⁴Cf. Azzouni [2004] for a separate critique of Quine’s criterion.

⁵Noneism is a version of Meinongianism, but whilst Meinong held that concrete objects exist, abstract objects subsist and possibilities and impossible objects do not exist in any way, Routley and Priest hold that concrete objects exist, and all other kinds of objects do not exist in any way; *viz.* they do not have an existential status weaker than but somehow analogous to existence.

⁶Azzouni [2004, p62] suggests this as a true sentence that quantifies over that which doesn’t exist.

This is to assume that the denizens of a world's domain are precisely the things that exist there. And this is rejected by noneism. If one is a noneist, there would seem to be no reason why the domain of each world should not be exactly the same, namely the set of all objects, whatever an object's existential status at that world. [Priest 2005: 13]

With the domain thus populated, sentences which we take intuitively to be correct such as 'There are fictional mice that talk'— $\exists x(Fx \wedge Mx \wedge Tx)$ —come out as true. But if one allows quantification over non-existent objects, then one risks trivializing the truth conditions for existentially quantified sentences. 'There are mice that have achieved grade 8 on the Stylophone' would be rendered true by the Stylophone-playing mice in D . The problem is this: if the truth of quantificational sentences is not constrained by what exists, then what constraints could there be? If the noneist wishes to retain this semantics then she may attempt a fix by building in more to the truth conditions of these sentences; only if there *exist* talking mice, *or* if (existent) people have pictorially represented talking mice, *or* if (existent) people have written about talking mice etc. is it true that some fictional mice talk. Unfortunately, once one specifies what properties and relations *existent* objects must have in order for sentences such as these to be true (assuming that such a procedure is possible), one will in doing so have produced what is in practice a Quinian regimentation of our sentences, rendering the noneist semantics otiose (at least for these purposes).

Moreover, the noneist provides a misleading picture of why we endorse sentences such as 'There are fictional mice that talk'. We assert 'There are fictional mice that talk' because of facts about existent objects; specifically, facts regarding cartoonists, animators and the like depicting talking mice. Correspondingly, we do not assert sentences such as 'There are fictional mice that play the stylophone' because there are no such facts concerning the depiction of stylophone playing mice. On the noneist picture however, these sentences are not correct because of the actions of the extant, but are made true by brute facts about non-existent objects.

4.3.2 A Quantifier for Metaphysics

A desideratum of a formal language is that it makes clear one's ontological commitments; i.e. it tells us what the world must be like in order for its sentences to be true. Gaining a grip on the ontological commitments of a theory Γ requires more than merely providing a list of the objects required to make Γ true; it involves also specifying the properties that those objects have, and the relations they stand in to other objects. If specifying ontological commitments was simply a case of listing objects, then revisionary understandings of the existential quantifier would be unproblematic in this respect; the truth conditions of a theory would be given by its models, while the ontological commitments of the theory would be specified by a subset \mathcal{D}_E of the domain \mathcal{D} of the model, where $\mathcal{D}_E = \{x \mid x \in \mathcal{D} \ \& \ x \text{ exists}\}$. The problem with this proposal is that it does not tell us the ontological commitments of Γ are, since it does not tell us what the properties of the objects in \mathcal{D}_E are. Nor can it, since, on this picture, some of

the objects in \mathcal{D}_E may stand in relations to non-existent objects. As such, we will not have a full list of the properties that extant objects must have in order for Γ to be true. The ontological commitments of Γ would have to be specified in a separate metalanguage. If the existential quantifier of the metalanguage does not express existence, then the same problem will arise; the ontological commitments of Γ will have to be specified in another metalinguage. If this language does not have an existential quantifier that expresses existence then the procedure will have to be carried out again, with no end in sight. This sequence will terminate only if appeal is made to an existential quantifier that expresses existence. If one accepts that formal languages ought to be intrinsically unambiguous with respect to their ontological commitments, then it follows that formal languages ought to have existential quantifiers that express ontological commitment.⁷

4.4 The Third Way

We now find ourselves in something of a dilemma. We have noted that speakers often quantify over objects that they do not mean to assert the existence of and treat the sentences conjoined in biconditionals such as (Num) as equivalent, yet we have argued that quantification expresses existence. Dropping the existential commitment of quantification was supposed to account for this linguistic data, but we have foreclosed the possibility of any such explanation. What is required is a means by which these two considerations—the non-existentially committal uses of quantification and the apparent equivalences expressed by (Num)—can be accommodated whilst retaining an ontologically committing existential quantifier. There is, however, a way to maintain all three commitments. The debate has been obscured by being framed in terms of the *semantics* of the existential quantifier: does ‘There exist numbers greater than 10^{10} that are prime’ *mean* that there *really are* numbers greater than 10^{10} which are prime? Does ‘There is someone who children think will bring them presents at Christmas’ *mean* that there *really is someone* who children think will bring them presents at Christmas? In fact, one can explain both the linguistic data marshaled by the revisionists *and* the the arguments in favour of the Quinian orthodoxy by shifting the focus away from semantics—what these sentences mean—to pragmatics—what speakers intend to convey by uttering these sentences.

⁷A caveat: there may be purposes for which it is useful to use a formal language which *is* ambiguous with respect to its truth conditions; we are free to create and use formal languages for whichever purposes we choose. The formal languages which ought to be unambiguous in this respect would be canonical in some sense; suitable for making unambiguous assertions about the world.

A second caveat: we must be open to the possibility that no language, in the ordinary sense, is suited to providing a fundamental description of the world. If ontic structural realism is correct, then formal languages quantifying, as they do, over domains of discrete objects will always misrepresent the world as containing such objects when in fact there are no things (on which various structures can supervene), but only structure itself. Having said this, a formal language supplemented with a function mapping the elements of its domain onto *positions* in structures may serve as a canonical language.

4.4.1 Two Ways of Thinking About Communication

The leading idea here is that, whilst the phrase ‘there are’ and its cognates express existence, speakers can *use* the phrase in such a way that what they convey or intend to convey does not involve existential commitment to (some or all of) the entities quantified over. Historically, views broadly of this sort have received short shrift in the philosophy of mathematics. John Burgess, in an influential article, tells us:

There is a *presumption* that people mean and believe what they say. It is, to be sure, a *defeasible* presumption, but some *evidence* is needed to defeat it. The *burden of proof* is on those who would suggest that people intend what they say only as a good yarn, to produce some actual evidence that this is indeed their intention. [Burgess 2004: 26]

and in a footnote:

To mean what one says literally is simply to mean what one says, just as to be a genuine antique is simply to be an antique. The force of ‘literally’ is not to assert that one is doing something more besides, but to deny that one is doing something else instead: meaning something *other* than what one says, as when one speaks metaphorically, hyperbolically, elliptically, or otherwise figuratively. One doesn’t have to think anything *extra* in order to speak literally: One has to think something extra in order to speak *non-literally*.

A skein of distinct claims are compressed into a short space here, and we would do well to untwine them. The claim that ‘people mean and believe what they say’ is related to and entailed by (although does not itself entail) a *Lockean* picture of communication: languages are codes—where codes can be thought of as systematic pairings of messages and signals—and communication takes place when one speaker firstly encodes their message into a signal (a series of sounds, inscriptions or gestures) which is heard or seen by another who subsequently decodes the signal and receives the message. The Lockean view stands in contrast to what we might call a *inferential* picture. The inferential picture treats communication as essentially involving abductive inferences about the intensions of fellow communicators based on their utterances and a number of contextual features. Interpretation in the Lockean picture merely involves decoding, whereas in the inferential picture, although decoding is involved, it is only one part of the interpretive process.⁸

Whilst few, if any, philosophers would *explicitly* endorse the Lockean picture as a theory of communication, it, or something very much like it, has conditioned the way that the debate over existential quantification has been carried out; and doubtless training in formal languages is wont to inculcate a certain picture of language and communication more generally.⁹ The Lockean picture is, however, cer-

⁸There is debate as to which parts of these inferences are personal or sub-personal and deliberate or spontaneous, and indeed whether (if they are unconscious and/or automatic) they deserve to be called ‘inferences’ at all. Nothing here hinges on how these debates are settled. It should be noted that I am using ‘inference’ in a very liberal sense, so as to include both spontaneous processes and those which take place at the sub-personal level. Those who dislike this liberal usage should feel free to make the requisite terminological adjustments.

⁹Stalnaker notes:

[P]ragmatics – the study of language in relation to the users of language – has been the neglected member of the traditional three-part division of the study of signs: syntax, semantics, pragmatics. The problems of pragmatics

tainly deficient as a theory of human *communication*. For one thing, natural language sentences are often ambiguous; viz. the signal is paired with two or more messages.¹⁰ Decoding then is insufficient to determine what a speaker intends to communicate. In addition, the assumption that people mean what they say does not do justice to the manner in which we *actually* communicate. Speakers often mean *more* than they say, for instance in the following dialogue:

Q: Do you like my new flatmate?

A: I think that anyone who talks about Eisenstein is pretentious.

We take the response to be an answer in the negative, but in doing so we go beyond the relatively parsimonious data encoded in the sentence and make an ampliative inference as to A's communicative intentions. Importantly, not only does what one conveys in uttering a sentence outstrip the semantic content of that sentence, commonly, on uttering a sentence, speakers fail to convey what is expressed by the sentence itself. One area in which this is particularly clear is the use of definite descriptions. Having just watched *Battleship Potemkin* with a group of friends, on uttering

The Eisenstein film was heavy-going.

one do not convey the content expressed by the sentence, i.e.:

[the x: x is an Eisenstein film] x was heavy-going

The content expressed by the sentence is false; there is more than one Eisenstein film. The sentence is uttered however in order to convey something that is true, i.e.:

[the x: x is an Eisenstein film & x = b] x was heavy-going (where the content of 'b' is the film b that this group of friends has just watched)

What the speaker here conveys does not entail what is expressed by the sentence she utters.¹¹ This phenomenon is widespread; as Kent Bach notes:

it is now a platitude that linguistic meaning generally underdetermines speaker meaning. That is, generally what a speaker means in uttering a sentence, even if the sentence is devoid of ambiguity, vagueness, or indexicality, goes beyond what the sentence means. [Bach 2005: 15-16]

Speakers will utter falsehoods when doing so is the most convenient way of conveying what they take to be true. People do not generally mean what they say. Nor do they have 'to think something extra in

have been treated informally by philosophers in the ordinary language tradition, and by some linguists, but logicians and philosophers of a formalistic frame of mind have generally ignored pragmatic problems, or else pushed them into semantics and syntax. [Stalnaker 1970: 272]

¹⁰Thus Groucho Marx: "Last night I shot an elephant in my pajamas, and how he got in my pajamas I'll never know."

¹¹Soames [2005, forthcoming] develops the above account of definite description use. See also Soames [2010] for a description of other ways in which the content of what speakers convey can come apart from the content of the sentences they utter. Soames intends his account as a normative one, but Sperber and Wilson [1986] present a similar inferentialist account of communication as an empirical hypothesis. See especially Sperber and Wilson [2008] for a description of how speakers can utter known falsehoods to convey what they take to be true.

order to speak *non-literally*'. Speakers, and their audience, are often not aware of the literal meaning of their utterance, only of what is communicated in making the utterance. Not only is this confirmed by our linguistic intuitions, there is empirical confirmation that literal interpretations of utterances are not given cognitive priority: non-literal and literal meanings are processed equally quickly; non-literal and literal meanings are processed in parallel, with neither having unconditional priority; and when both literal and non-literal meanings are available (i.e. are relevant in the context of utterance), both are processed.¹² In cases such as the definite description use described above, speakers are not in general aware of the content expressed by the sentence they utter, *viz.* that there is exactly one Eisenstein film and that it is heavy-going.¹³ Nor is it true that when speakers do not mean what they say, they intend their utterance as a 'good yarn': a person who utters 'The Eisenstein film was heavy-going' is *not* merely story-telling. The exclusive dichotomy between literalness and not attempting to express anything about the world, is a false one.

Once one recognizes this, one opens up space for a picture of quantificational language that can accommodate the linguistic data—i.e. the non-existential *use* of the 'there is' idiom—whilst retaining a standard semantics for quantification: 'There are' and ' $\exists x$ ' *do* express existence, but speakers can *make use* of quantificational sentences to *convey* beliefs that do not. This picture reconciles what had appeared to be conflicting considerations, and it also offers a diagnosis of Priest and Azzouni's conclusions. It is well known that we are prone to inappropriately import pragmatic intuitions into our semantic evaluations of sentences. In particular, it is easy to conflate the meaning of a sentence with what that an utterance or inscription of that sentence would be used to convey in ordinary contexts. For example, if a speaker utters "I tried to climb Ben Nevis" whilst describing a successful expedition to the top of the mountain, there is a strong intuition that the sentence is not true. But once we realise the source of this intuition (that the utterance is pragmatically inappropriate) we realise why it is misleading. Conversely, when a sentence is pragmatically appropriate, we are inclined to assess it as being true, even when it is not (as in the use of definite descriptions).

A similar slip is made by primitive expressivists, who move from the observation that people utter sentences such as "Murder is wrong" in order to convey disapproval of murder to the conclusion that the *meaning* of the sentence 'Murder is wrong' is to be equated with those pragmatic effects. In doing so they commit the pragmatic fallacy [cf. Kalderon 2005, ch.2]. Those who conclude that the existential quantifier does not express existence on the basis of the way we use quantificational sentences in the vernacular make an equivalent mistake. Here, the Meinongian¹⁴ is mistaken about what her intuitions are tracking. She is attending to what information she would ordinarily use such sentences to convey and mistakenly importing those intuitions onto what the sentence *means*. There is however a temptation to

¹²Cf. Glucksberg [2008] for a summary of the relevant empirical data.

¹³This is not restricted to definite description use. Recanati [2004, ch1] provides an overview of the oft-discussed examples: (1) I've had breakfast; (2) You are not going to die; (3) It's raining; (4) The table is covered with books; (5) Everyone went to Paris; (6) John has three children.

¹⁴It should be noted that Azzouni does not consider himself a Meinongian, and Priest does not endorse the Meinongian tenet that abstract objects *subsist*, instead taking them to not exist in any sense. Here I'm using the term (perhaps unfairly) to describe anyone who does not take the existential quantifier to express existence.

commit an equal but converse *semantic* fallacy by moving from the observation that ‘there are’ sentences express existence to the conclusion that people *use* ‘there are’ sentences in order to convey existence. But not only is there no reason to assume that this is so; the assumption is incongruous with the facts of language use. So, for example, someone who utters “There are three dents in my car” does not necessarily believe that there *are* dents, nor wishes to express the belief that there are dents. However, they can use the sentence to convey beliefs about the topographical properties of their car.¹⁵ Not only does this understanding of quantification dissolve the standoff between Quinians and Meinongians by unearthing what is right about each; it explains why philosophers who insist that they find noneist quantification unintelligible aren’t simply being obtuse in the face of linguistic data which seemingly contravenes their doctrine—they are attending purely to the semantics of the existential quantifier.

Why opt for a pragmatic explanation rather than taking ‘there is’ to be ontologically neutral¹⁶ or ambiguous, requiring contextual factors to provide disambiguation? Treating the existential quantifier as ontologically neutral was intended to accommodate the linguistic data. However, ontologically neutral quantifiers have difficulty accommodating linguistic data themselves. It is often noted—for instance by Azzouni [2004] himself—that the use of emphasis, either by way of emphatic intonation or by the introduction of worlds such as ‘really’, is employed to indicate when ontological commitment is intended in existentially quantified utterances: e.g. “There *really are* numbers.” This practice however, does not make sense if one takes ‘there is’ to be ontologically neutral or ambiguous, since the use of intonational emphasis or ‘really’ do not act either to themselves express existence or to disambiguate a term, as would be required if the existential quantifier were either ontologically neutral or ambiguous (consider: ‘I really went to the Bank’). Rather, they function to indicate that the speaker intends her utterance to be interpreted literally. Their effect on the ontological commitment of existentially quantified sentences is then predicted by the pragmatic account, but precluded by Priest and Azzouni’s ontologically neutral quantifier.

4.4.2 Embellished Worlds

‘There is someone who children believe will bring them presents at Christmas’, ‘There are three dents in my car’ and the like really are informative, but according to the quantificational Quinian they also really are *false*. In order for the Quinian to explain these uses of the existential quantifier what is required is some account of how quantification over things that do not exist can be informative with respect to that which does not exist. What is wanted then is some notion of *adequacy* with respect to the extant: a way in which one can understand how strictly false sentences can ‘get things right’ with respect to what exists. Intuitively, a claim is adequate with respect to the extant, when existent objects are the way they would have to be in order for the claim to be true.¹⁷ As a means to flesh out this notion somewhat,

¹⁵Consider also Azzouni’s quantificationally loaded paragraph (*).

¹⁶As Azzouni argues (cf. esp. [2007]).

¹⁷What I am calling ‘adequate’ Leng calls ‘fictional’. See Leng [2010: ch8] for an overview of ways of characterizing nominalistic adequacy or *fictionality*.

I will here outline a semi-formal framework. Central to this framework is the idea of an *embellished world*. An embellished world is a pair $\langle w, e \rangle$ containing a world w (with a built-in domain D_w) and an *embellishment* e : a domain D_e which supplements that of w , such that the domain $D_{w \cup e}$ of $\langle w, e \rangle$ is $D_w \cup D_e$. We can describe embellished worlds using a 2-sorted language¹⁸ $\mathcal{L}(\langle I_i \rangle, \langle M_i \rangle, \langle E_i \rangle)$ which ranges over:

- (1) *intrinsic entities* (objects in the world), using primary variables x_1, x_2, \dots, x_n .
- (2) *extrinsic entities* (objects in the embellishment), using secondary variables y_1, y_2, \dots, y_n .

As such the language will contain three kinds of predicate:

- (i) *intrinsic predicates*, expressing relations between objects in the world: $\langle I_i \rangle = (I_1, I_2, \dots)$
- (ii) *extrinsic predicates*, expressing relations between objects in the embellishment: $\langle E_i \rangle = (E_1, E_2, \dots)$
- (iii) *mixed predicates*, expressing relations between intrinsic and extrinsic objects: $\langle M_i \rangle = (M_1, M_2, \dots)$

We will say that a sentence of the mixed language is *correct* iff_{df} it is true in the embellished world. Correct sentences can be false, but they are always *veridical* with respect to D_w . More specifically, sentences are correct or *adequate* with respect to a certain ontology. Let us say that w is wholly concrete; in this case a mixed mathematico-physical sentence is correct or nominalistically adequate iff it is true in $\langle w, e \rangle$, where e is an embellishment with the requisite mathematical objects. Let us say that the w is entirely composed of mereological simples; then a sentence quantifying over macroscopic objects is ‘nihilistically’ correct in w if it is true in $\langle w, e \rangle$: a world composed of simples, embellished with the requisite macroscopic objects.¹⁹ Utterances that are adequate, are a means of describing the world indirectly, by placing restrictions on the way *extant* objects are.

Linguistic Data

In the philosophy of mathematics, two reasons for doubt arose for the existential commitment of \exists ; one was the non-existential use of ‘there is’ in the vernacular, and the other was the seemingly trivial derivations of substantive ontological theses. The pragmatic account, making use of embellished domains, can provide a unified treatment of both phenomena. The ‘there is’ idiom is frequently used by speakers who do not intend to express existence; as by quantifying over objects she does not believe in, the speaker can convey information about the objects she *does* believe in. For instance, if a nominalist utters “The

¹⁸Burgess and Rosen [1997] and Ketland [2004] adopt a similar framework to discuss mixed abstract/concrete languages and mixed theoretical/observable languages respectively.

¹⁹Rosen [2001] defines nominalistic adequacy by appealing to the ‘concrete core’ of a world: a sentence S is nominalistically adequate iff ‘the concrete core of the actual world is an exact intrinsic duplicate of the concrete core of some world at which S is true’ [Rosen 2001: 75]. Ketland [2011] defines nominalistic adequacy in model-theoretic terms. Our language \mathcal{L} has an intended interpretation I , where I° is the concrete part of I . An \mathcal{L} -structure is a nonempty domain U and an interpretation assigning an element of U to each constant in \mathcal{L} , a function from U^n to U to each n -ary function in \mathcal{L} and a subset of U^n to each n -ary relation in \mathcal{L} . An \mathcal{L} -structure is nominalistically correct iff the concrete part of it is isomorphic to I° . Defining adequacy in the way we have done serves to accent the relationship between adequacy and logical consequence; a matter we will turn our attention to in section 5.

number of divine persons is three” the sentence expresses a relation holding between intrinsic entities—the divine persons—and extrinsic entities—the positive integers. This is nominalistically correct so long as it is true at the actual (concrete) world embellished with the integers.²⁰ Here, we can articulate directly what needs to be *true* at the actual world for the sentence to be nominalistically correct; i.e. that there are three divine persons, but this need not always be the case. For instance, if a nominalist utters “The mass in kilograms of d_1 is $\sqrt{2}$ the mass in kilograms of d_2 ” this will be nominalistically correct iff $\exists y_1 \exists y_2$ in the embellished world such that $(y_1 = \text{Mass}_{kg}(d_1) \ \& \ y_2 = \text{Mass}_{kg}(d_2) \ \& \ y_1 = \sqrt{2}y_2)$. In this case it is far more difficult to give an intrinsic description of what obtains at the actual world when this is nominalistically correct. But herein lies an important motivation for this pragmatic use of sentences; sometimes it is simply too difficult or time consuming to describe features of the world directly, but we can do so *indirectly* by imagining our world to be embellished in certain ways. By uttering a mixed sentence M , the nominalist conveys the information that concrete objects have the properties they would have to possess in order for M to be true.²¹ It is beyond the capacities of many, if not most, speakers to describe the topological properties of their car directly, but by embellishing the world with dents they can convey the requisite information more easily.

Metametaphysics

The pragmatic approach also suggests an explanation of our tendency to treat some nominalistic sentences as equivalent to more ontologically loaded counterparts, whilst retaining an existentially committing $\exists x$. Sentences such as those conjoined by a biconditional in (Num) are—even given a Fregean interpretation of (Num)—*pragmatically* equivalent; i.e. they are assertable in the same contexts.²² So long as the speaker’s goal is to utter sentences which are nominalistically correct (rather than literally true) then they will *treat* nominalistic sentences and their abstract counterparts as equivalent. Seen in this light, we have defused at least one motivation for metaphysical anti-realism. Our tendency to speak and make inferences as if we are ignoring the addition of new entities into our domain need not be seen as reflection of the insubstantiality of these entities (and hence an indication that the existential quantifier does not express ‘robust’ existence), but merely as a practical strategy for communication. In addition to the biconditional (Num), this allows for a highly deflationary view of abstraction principles, such as ‘Hume’s Principle’:

²⁰Or, for the nominalist, *would be true were* the actual world embellished with the integers.

²¹Of course, this does not only take place in the case of *mathematical* abstracta; other kinds of abstract talk can be used to implicate facts about concreta. If an ordinary speaker utters “There’s no such thing as fascism”, they mean to convey that there are no fascists. Note however that there is no semantic contradiction in uttering “There’s no such thing as fascism, but there are many fascists”. The incompatibility between ‘There’s no such thing as fascism’ and ‘There are many fascists’ arises because linguistic communities tacitly accept bridge principles associating fascism and fascists. As with fascism, so with fiction. ‘Harry Potter is more successful than all real wizards’ is *correct* because linguistic communities accept (as correct) bridge principles linking the creative activity of authors with the fictional characters they describe. Platonists who attempt to capture the correctness of this Harry Potter description in terms of *truly* describing abstract objects (such as van Inwagen [2000]) are lumbered with the consequence that the content of novels is true. I take it that it is an advantage of the current account that it treats fictions—whether they be Harry Potter or The Protocols of the Elders of Zion—as *false*.

²²At least, they are assertible in *most* of the same contexts. In some philosophical contexts where the different meanings of each side of (Num) are *salient* they are not assertible in the same contexts. That they *are* assertible in the same contexts in everyday settings appears to provide some evidence that speakers (ordinarily) aim for nominalistic adequacy rather than truth *tout court*.

The number of F s = the number of G s iff there is a one-to-one correspondence between the F s and the G s.

which has the form:

$$\forall\alpha\forall\beta(\Sigma(\alpha) = \Sigma(\beta) \leftrightarrow \alpha \approx \beta)$$

where Σ is an appropriate term-forming operator and \approx an equivalence relation. We can explain the strong intuition that there is something right about abstraction principles in terms of their a priori nominalistic adequacy and their attendant assertability in all contexts, without having to make opaque appeals to ‘insubstantial’ facts.

In trying to understand the existential quantifier we removed one reason for adopting metaphysical anti-realism, but in doing so we may also have removed a barrier to a workable nominalism. Carnap famously complained that the nominalist who helps herself to a vocabulary of abstract terms whilst denying the existence of abstract objects in the seminar room was ‘like a man who in his everyday life does with qualms many things which are not in accord with the high moral principles he professes on Sundays.’ [Carnap 1950: 20]. Thirty years on Field [1980: 2] complained that the fictionalist who takes back in his philosophical moments what he asserts whilst doing science is guilty of ‘intellectual doublethink’. Field’s brilliant but quixotic response was an attempt to purge science of its abstract vocabulary, but here a different route suggests itself. If we do not assert sentences about abstracta, but utter them with an eye to nominalistic adequacy then we can *use* platonistic language without committing ourselves to its concomitant ontology.²³

4.5 Applications

Although debates over mathematical nominalism have tended to focus on the applicability of mathematics as it pertains to physics, there is a less well discussed, but equally pressing, issue confronting nominalists regarding the relationship between applied mathematics and logical consequence. As we will see, one application of our framework consists in providing a nominalist account of the so-called ‘speed-up’ phenomenon, which appears problematic for nominalism. Speed-up is said to take place when a sentence ϕ can, in principle, be derived from one system, but where a far shorter derivation is available in a system with stronger existence assumptions. The cases perplexing to the nominalist are those where some *concrete* body of assumptions N is sufficient to derive some ϕ in the language of N , but where (i) this proof would be extortionately large, far too large to ever be carried out in practice,

²³A proviso: While I am sympathetic to the thought that what we have presented may represent an accurate, if highly idealized and speculative, picture of the communicative processes of most actual speakers (who, outside the philosopher’s study, are usually unconcerned about the existence of dents whilst deeply concerned about the state of their cars), we should think of this picture not as armchair linguistics or psychology, but as a rational reconstruction of the nominalist’s use of quantification over abstract objects. Perhaps some or most nominalists really are engaged in intellectual doublethink when they casually quantify over abstracta, but they *need not* be.

and (ii) where a far shorter derivation is available from a mathematical theory $S^{24} + N$. Ketland [2005, p20] provides an example of this with the following inference I_1 :

- (1) The number of people in the room is 100.
- (2) The number of houses in the street is 99.
- (3) Each person in the room lives in exactly one house in the street, therefore
- (4) At least two people in the room share the same house.

(1)–(4) can be expressed nominalistically in the familiar way, by formalizing ‘There are exactly n F s’ as ‘ $\exists x_1, \dots, \exists x_n (x_1 \neq x_2 \wedge \dots \wedge x_{n-1} \neq x_n \wedge (Fx_1 \wedge \dots \wedge Fx_n \wedge \forall y (Fy \rightarrow y = x_1 \vee \dots \vee y = x_n)))$ ’, and the inference I_1 is valid in first-order logic. Unfortunately, Ketland estimates that the shortest derivation of (4) from (1)–(3) would be more than 10^8 symbols in length.

If we make use of some mathematical assumptions however, (4) can be derived much more easily. In particular, if S includes ZFC and the axiom of urelements, then (4) can be derived in relatively few lines. The ‘Pigeonhole Principle’ says that for sets X and Y , if $|X| > |Y|$ then any function $f : X \rightarrow Y$ will be non-injective. This is true by the definition of $>$ in set theory. Letting X be the set of people in the room and Y the set of houses in the street it can be quickly proven that at least two people in the room share one house. Ketland lays down the following challenge for nominalists:

presumably the nominalist does not wish to deny the validity of the inferences ... under consideration. But there is no feasible direct verification for the above inferences, and the short mathematical derivations involve *practically indispensable* assumptions about numbers, sets and functions. So, how might a nominalist account for our knowledge that such inferences are valid? [Ketland 2005, p23]

Ketland states this as a challenge, but we can put more pressure on the nominalist by turning it into an argument: I take it as undeniable that we *do* in fact know that (4) can be derived from (1)–(3). Moreover, it is wholly unfeasible that we know (4) on the basis of a deduction from N . Given that we can’t know that (4) can be derived from (1)–(3) on the basis of N , *if* we know that (4) can be derived from (1)–(3) then we must know the mathematical assumptions S . It follows by *modus ponens* that we must know S . This implies that nominalism is false. The nominalist must then deny the conditional, but how?

One line of response the nominalist may be inclined to take is to argue that we know (4) on some non-inferential basis; perhaps we can visualize the people and houses and *just see* that at least two people must share the same house. This is not implausible, however there is no reason to think that the strategy could be made to apply in all cases; and indeed there are cases where it seems clear that it could not (Boolos [1987] provides one such example, also discussed by Ketland [2005]).

²⁴ S will not merely consist of pure mathematical theory; if it were, the separate vocabularies of S and N would trivially ensure that adding S to N would not allow us to make any additional inferences to conclusions in the language of N (provided S is consistent). S must contain *bridge laws* that involve both mathematical and non-mathematical vocabulary in order to yield new inferences.

Here I think that our apparatus of embellished worlds can show how the nominalist *can* meet this challenge. We said that a sentence is ‘correct’ or nominalistically adequate iff_{df} it is true at its associated embellished world. Although $S + N$ is false (according to the nominalist) it *is* nominalistically adequate. Of course, truth is preserved under deduction, but so too is nominalistic adequacy. This follows from our definition of nominalistic adequacy as truth at an embellished world; if Σ is true at $\langle w, e \rangle$ and $\Sigma \vdash \phi$ then ϕ is true at $\langle w, e \rangle$. Accordingly, if the nominalist knows that $S + N$ is nominalistically adequate and that $S + N \vdash \phi$ then she thereby knows that ϕ is nominalistically adequate. But if some ϕ is nominalistically adequate and only quantifies over concrete objects then it is true; this much follows from platitudes about how to evaluate quantified sentences in a model. In this way the nominalist can come to know the *truth* of some ϕ from the *nominalistic adequacy* of some Σ . I will call this process *bridging*; as we can use sentences we take to be false to bridge deductions from N to ϕ .

So far, we have shown that the nominalist can know the truth of (4) from the nominalistic adequacy of (1)-(3) + S , but our goal is to understand how the nominalist can know that (4) follows from the truth of (1)-(3) *alone*. The platonist is in a similar boat: we have seen that they can know the truth of (4) from the truth of (1)-(3) + S ; but the platonist must also account for how it can be known that (4) follows from the truth of (1)-(3) *alone*. In order to do so we must be warranted in holding that S places no constrictions on the way things are with regards to contingent physical facts. For the platonist, this will amount to S being necessarily true; for the nominalist, this will amount to S being conservative.²⁵ Not all applied mathematics is conservative: bridge laws may encode information about the concrete domain that does not hold of necessity. In this case there exists a proof that the applied mathematics *is* conservative [Field 1980]. It may be thought that the platonist has an advantage over the nominalist here. The proof, after all, is a mathematical proof, and the platonist may claim that the nominalist has no right to appeal to it. Here the nominalist has an *ad hominem* retort. Nominalists and platonists alike should, on the basis of the proof, accept the conditional ‘If platonism is true then S is conservative over N ’, which is equivalent to ‘Either platonism is false or S is conservative over N ’. If the first disjunct is accepted then the nominalist is home and dry. If the second is accepted then the nominalist may make use of the conclusion of the proof.²⁶ However, the wider point here is that there is no epistemological problem peculiar to the nominalist: because of the close association between necessary truth and conservativeness, in those cases where the platonist is warranted in holding that a particular inference is valid, the nominalist will be likewise warranted; for whenever the platonist is warranted in taking S to be necessary, the nominalist is likewise warranted in taking S to be conservative.²⁷

²⁵Note however that there is no reason for the nominalist to hold that *all* applied mathematics must be conservative, unless they hold that all mathematics is theoretically dispensable: bridging can take place whenever $S + N$ is nominalistically adequate, regardless of whether it is conservative over N .

²⁶The point is quite general: for this reason nominalists can make use of any platonistic proof in the defence of nominalism.

²⁷A note on Field’s conservativeness proof: If N contains an axiom scheme then applied set theory is only guaranteed to be conservative if the axiom scheme is stated as a list (of the form $\neg Q_0, \neg Q_1, \neg Q_2, \dots$). If the axiom scheme is stated as a rule (for every formula Q , $\neg Q$ is an axiom), then applied set theory is non-conservative over N . Cf. Burgess and Rosen [1997, pp.194-6] for discussion.

The platonist may here feel that some sly nominalist casuistry is at hand: the nominalist should not be allowed to help herself to the benefits of mathematics without accepting its supposed ontology. I think we can mitigate this feeling, although perhaps not to the platonist's full satisfaction, by considering how it may have arisen. One concern may be that the bridging phenomenon is really just a form of if-thenism in disguise. If-thenism is often stated as the view that what mathematical statements, such as '71 is prime', really *mean* is some conditional statement such as 'If numbers exist then 71 is prime'. Dropping this hermeneutic element we could endorse a form of if-thenism in which we only believed conditional versions of mathematical claims. Instead thinking that ϕ is true on the grounds that Σ and that $\Sigma \vdash \phi$ are true we could merely endorse $\Sigma \rightarrow \phi$. The problem with this approach is that it does not give us unconditional knowledge that ϕ , and our unconditional knowledge that ϕ is precisely what requires explanation in these cases. However, with our apparatus of embellished worlds we are able to say something *unconditional*—that Σ is nominalistically adequate—and it is unconditional knowledge of *this* that allows us to infer ϕ . Another concern may be that the nominalist is illicitly helping herself to mathematical assumptions. We have said that Σ is nominalistically adequate iff it is true at $\langle w, e \rangle$; but $\langle w, e \rangle$ is an abstract object and so our definition of nominalistic adequacy itself is inconsistent with nominalism. This much is fair, however there is a simple fix for the nominalist. The nominalist should endorse:

1. Σ is nominalistically adequate \leftrightarrow_{df} were w embellished with e then Σ would be true.
2. If Σ is true at $\langle w, e \rangle$ and $\Sigma \vdash \phi$ then ϕ is true at $\langle w, e \rangle$ (i.e. if Σ is nominalistically adequate and $\Sigma \vdash \phi$ then ϕ is nominalistically adequate).
3. If ϕ is nominalistically adequate and ϕ only quantifies over concrete objects, then ϕ is true.

and in doing so can still infer the truth of ϕ from the nominalistic adequacy of Σ but without assuming the existence of embellished worlds.²⁸

4.5.1 Losing Our Grip on the World

I want to briefly address an objection that is commonly raised to this sort of approach. An inability to describe the world directly, using straightforwardly true sentences amounts to a kind of *mysticism*; a descriptive or epistemic scepticism. If one cannot describe the world directly, using straightforwardly true sentences, then one has no grip on what the world is like. We will address this issue in far greater detail in chapter 5, but for the time being it would be gainful to say something to assuage this fear. Consider the nominalistically adequate sentence 'There are three dents in my car'. I have claimed that, despite its

²⁸It may be objected that this requires modal knowledge on the part of the nominalist, and that *this* involves a commitment to abstract entities. I disagree. A proper defence of this would take us too far afield, but, with Field [1989, 1991] I believe there is good reason to think that modal notions are primitive (see also Leng [2007] for a useful elaboration and defence of this). It is worth noting however, that the modal knowledge (if indeed it deserves the appellation 'modal knowledge') involved here is *very* modest indeed; it amounts to nothing more than knowing how to assess truth in a domain; something that is just as available to the nominalist as to the platonist.

falsehood, this carries real information about the world. Those who deny that merely nominalistically adequate sentences can do such a thing face a dilemma. On the one hand, if they deny that utterances of the sentence are genuinely informative, then they make a mystery of our linguistic practices; for linguistic communities *do* make use of utterances of this sort, and manage to convey information by doing so. If, on the other hand, they accept that merely nominalistically adequate utterances can be used in this way, whilst retaining the commitment that false sentences cannot be used to describe the world, then they must hold that what speakers *really* grasp when they hear ‘There are three dents in my car’ is some nominalization of the sentence, that does not quantify over dents but describes the topological properties of the car directly. This however is psychologically implausible. Nominalization of this sort, is complicated and requires specific training; it is not part of the quotidian linguistic practices of ordinary speakers.

4.6 Summation

We have provided an outline of how to understand existential quantification. It is only an outline, but it is an outline that suggests a route out of a philosophical impasse; perhaps the *only* route that can do justice to, on the one hand, our semantics and the need for a useful existential quantifier to express existence, and on the other hand the facts of how we communicate and our tendency to treat ontologically innocent statements as equivalent to their ontologically loaded counterparts. We subsequently turned our attention to logical inference and the speed-up phenomena. Herein, I think, lies the power of our approach. One instance of the speed-up phenomena served as a test case, but there is scope to extend the framework to shed light on the relationship between abstract ontology and inference more generally, particularly with regards to the correlation between stronger existence assumptions and deductive power. A particular area of interest is that of higher-order logic, where it is generally held that the additional deductive power afforded by the logic comes at the price of a set-theoretic ontology. Boolos [1984] showed how to construe higher-order logic in an ontologically innocent way, by treating second-order quantification as plural generalization. The ontological gains are however counterbalanced by a coincident loss in our expressive and inferential resources; for instance in understanding quantification over relations. Our considerations are suggestive of a different avenue of possibility: one where the guiding idea is *adequacy* with respect to an accepted ontology; one which would allow us to have our cake and eat it too.

Chapter 5

Mathematics and the World

5.1 The Indispensability Argument

In this chapter we will address what I take to be the primary challenges towards the mathematical nominalist. In the course of the chapter a number of different topics will be addressed: the relationship of nominalism to realism—whether nominalism without a successful programme of nominalization entails, in Burgess and Rosen’s words, ‘that standard science and mathematics are no reliable guides to what there is’; the nature of representation, and, in particular, scientific and mathematical representation; the nature of explanation; the relationship of pragmatism to these issues; and Quine’s attempt to produce a *method* of determining ontological commitment. But our route through these issues will be via *indispensability arguments*.

Although the indispensability argument is customarily attributed to Quine and Putnam¹, versions of the argument are sometimes traced back to Frege [1970, i.e. Grundgesetze Volume II, §91] and Gödel [1947]². In a way it is testimony to the ubiquity of the argument in the minds of philosophers that it is rarely expounded in a systematic way. I take my formulation of the argument from Patrick Dieveney (with minor modifications) [Dieveney 2007]³:

1. *The Indispensability Thesis*: Our best scientific theories indispensably use mathematical principles and indispensably refer to mathematical objects, i.e., the sentences resulting from formalizing these theories in predicate logic indispensably include existential quantification over mathematical objects.
2. *Confirmation Holism*: Evidence for a scientific theory confers justification on the whole theory (including its mathematical component) and not on the individual components of the theory sep-

¹C.f., for example, Putnam [1971], Quine [1953] and a number of later writings [such as 1981].

²Pieranna Garavaso [2005] argues that the attribution of the indispensability argument to Frege is mistaken.

³Other systematic expositions of the argument can be found in Maddy [1997], Resnik [1997], Colyvan [2001] and Leng [2010].

arately.

3. *Theory Realism / Naturalism*: We look to our best scientific theories as the ultimate arbiter of existence and truth.
4. *Quine's Criterion*: The ontology that acceptance of a theory commits one to is determined by what objects are required to satisfy the existential consequences of that theory.

The charge leveled against the nominalist is that she makes *use* of scientific theories that quantify over mathematical objects, but does not accept the ontological commitments which these uses give rise to. Her behaviour is at odds with her beliefs; she is engaged in 'intellectual doublethink'. Hartry Field puts it like this:

An indispensability argument is an argument that we should believe a certain claim (for instance, a claim asserting the existence of a certain kind of entity) because doing so is indispensable for certain *purposes* (which the argument then details). [Field 1989: 14, my emphasis]

Notice that indispensability arguments, as Field understands them here, involve extracting ontological commitments from *practical* commitments. Quantification over such and such an entity is indispensable for certain *purposes*; there are things that cannot be *done* without quantification over mathematical objects. Hence, doing those things involves a commitment to mathematical objects. Field of course contests the first premiss of this argument. He argues that our best scientific theories need not quantify over mathematical objects at all—i.e., that science can be nominalized—and so that the things one can do with mathematical science can also be done with nominalistic science. The approach taken here will be very different. Rather than contending that one can do with nominalistic language what one can do with mathematical language, my focus will be on the association between use and commitment *itself*. Our road into that issue however will be through the indispensability argument.

5.1.1 Premiss 4

Premiss 4 is, as it is stated, ambiguous. There are different ways in which one can 'accept' a theory, for instance as empirically adequate (van Fraassen [1980]), nominalistically adequate, or straightforwardly true, and no doubt there are other flora and fauna of acceptance besides these. To accept a theory as true however, we must accept the existence of the objects that theory says exist, and as I argued earlier, quantification expresses existence. So far we are in agreement with the advocate of the indispensability argument.

5.1.2 Premiss 2

The status of the second premiss is somewhat controversial. In fact, it is not obviously required for the indispensability argument, but is sometimes thought to provide the grounds for accepting premiss 3. For this reason Mark Colyvan gives confirmation holism a central role in his defence of platonism:

Confirmation holism is the view that theories are confirmed or disconfirmed as wholes. So if a theory is *confirmed* by empirical findings, the *whole* theory is confirmed. In particular, whatever mathematics is made use of in the theory is also confirmed. Furthermore ... the same evidence that is appealed to in justifying belief in the mathematical components of the theory is appealed to in justifying the empirical portion of the theory' [Colyvan 2001: 13]

and

According to the holist, mathematical beliefs are justified in exactly the same way as other beliefs: by their role in our best scientific theories and these, in turn, are justified by appeal to the usual criteria of theory choice (empirical adequacy, simplicity, explanatory power, and so on). [Colyvan 2007: 111, my emphasis].

As a result of this, some nominalists have attacked the indispensability argument by attacking confirmation holism. As Mary Leng puts it in her recent book:

The mere presence, in our best presentations of our scientific theories, of sentences whose literal truth would require the existence of mathematical objects, will not give us reason to believe that there are such objects, if those sentences are not amongst the parts of our theories whose literal truth should be considered confirmed by our ordinary scientific standards. [Leng 2010: 101]

and

We need to look more closely at the question of *which parts* of our scientific theories receive confirmation from our theoretical successes, in order to discover whether we ever have reason to believe our theoretical claims whose truth would require the existence of mathematical objects.' [Leng 2010: 131, my emphasis]

In the same vein, but with a different goal from either Colyvan or Leng in mind, Eliot Sober [1993] argues that indispensability arguments *fail* precisely because confirmation holism is incompatible with what he takes to be the most plausible theory of confirmation: *contrastive empiricism*. Contrastive empiricism takes as central to our understanding of confirmation a *Likelihood Principle*, given by Anthony Edwards [1972]:

Observation O favours hypothesis H_1 over H_2 if and only if $P(O/H_1) > P(O/H_2)$.

One feature of the Likelihood Principle is that it entails a symmetry between confirmation and disconfirmation; an observation O confirms a hypothesis H if and only if the absence of O disconfirms H .

$$P(O/H_1) > P(O/H_2) \text{ iff } P(\neg O/H_1) < P(\neg O/H_2)$$

We are to suppose that each of our competing empirical hypotheses (H_1, H_2, \dots, H_n) contain a set M of mathematical propositions, and that O favours one of these hypotheses over the others. Since M is a part of every hypothesis however, O cannot be said to favour M over any other hypothesis, and so M is not confirmed empirically but is a ‘background assumption common to the hypotheses under test’ [Sober 1993: 45].⁴

There is I think a danger inherent in couching the debate in terms of confirmation holism. That confirmation holism could play such a central role seems to be ruled out by the logical structure of mathematicized theories. In order to make out a notion of nominalistic adequacy, we adopted earlier a two-sorted language $\mathcal{L}(\langle I_i \rangle, \langle M_i \rangle, \langle E_i \rangle)$. Similarly, to describe mixed mathematical-physical theories, one can adopt a two-sorted language $\mathcal{L}(\langle C_i \rangle, \langle M_i \rangle, \langle A_i \rangle)$ which ranges over:

- (1) *concrete entities*, using primary variables x_1, x_2, \dots, x_n .
- (2) *abstract entities* (objects in the embellishment), using secondary variables y_1, y_2, \dots, y_n .

and contains three kinds of predicate:

- (i) *concrete predicates*, expressing relations between concreta: $\langle C_i \rangle = (C_1, C_2, \dots)$
- (ii) *abstract predicates*, expressing relations between abstracta: $\langle A_i \rangle = (A_1, A_2, \dots)$
- (iii) *mixed predicates*, expressing relations between concrete and abstract objects: $\langle M_i \rangle = (M_1, M_2, \dots)$

The difficulty with framing the debate in terms of confirmation holism arises because the kind of things that can be, or can fail to be, confirmed are things that express propositions. As such whatever parts of theories we take to be confirmed must come in “proposition-sized” chunks, *viz.* sentences. It does not, for instance, make sense to say that a *thing* could be confirmed, whether that thing is a function from the integers to the reals, or a tin of paint. *That* a function from the integers to the reals, or that a tin of paint exists, on the other hand, are at least candidates for confirmation. This matters because mathematicized theories cannot in general be neatly sundered into nominalistic and mathematical parts. If such a thing were possible, the nominalist could, without doing damage to the picture of science as a guide to the nature of the world, simply argue that the nominalistic part was confirmed whilst the mathematical part was not, and leave the matter there. Mathematicized theories are however, in large part, a kind of abstract-concrete alloy. As is exhibited when one makes clear the logical structure of two-sorted languages, whilst one can divide the *language* of mathematicized theories into a nominalistic part—ranging over concreta and expressing relations (including one-place relations) between these objects—and a mathematical part, ranging over abstracta and expressing relations (including one-place relations) between those objects—the *theory itself will consist of three kinds of sentence*: those ex-

⁴How much comfort the nominalist is to take from this conclusion is not exactly clear. Even if the existence of mathematical objects is not confirmed empirically, if assuming their existence is a prerequisite for carrying out scientific study then the nominalist is not better off. What this shows is that there are stronger and weaker ways of understanding the indispensability argument. The stronger hopes to show that the existence of mathematical objects is confirmed by observation; the weaker that the existence of mathematical objects must be presupposed in order to coherently practice science. Even the weaker claim however is incompatible with nominalism.

pressing relations between concreta, those expressing relations between abstracta, and those expressing relations between concreta and abstracta. The difficulty with framing the nominalist's task in terms of confirmation holism is that many of the sentences of a mathematicized theory that are of interest are mixed sentences. Without mixed sentences mathematics would not be applicable (or, at least, we would have no way of describing how it is applicable); for bridge-laws are, of necessity, mixed sentences expressing relations between mathematical objects and concrete ones. Absent bridge laws mathematicized theories would consist of a pool of mathematical statements, a pool of nominalistic statements, and no instructions as to how the mathematics can be *applied* to the nominalistic subject matter. In addition, much of what a mathematicized theory says about the concrete domain is expressed in mixed sentences. Measurements, for instance, are expressed in terms of functions from concreta to real numbers; we describe physical quantities by associating them with numerical magnitudes. 'The mass of d_1 is 5 kilogrammes' would be formalized as $5 = \text{Mass}_{kg}(d_1)$, which entails $\exists x(x = \text{Mass}_{kg}(d_1))$. If only the nominalistic part of mathematicized theories are confirmed, then while we may be able to assert the existence of the requisite concreta, we will not be able to describe them in any interesting way. Often concreta don't even get a mention *qua* concreta. Physical systems are represented *as* mathematical structures; concrete objects are reconstituted in the theory as abstract objects and implications about *these* objects are deduced. Planets are described as point masses, gasses as collections of extensionless molecules moving at uniform velocities on a two-dimensional plane. Even observation sentences may be couched in mathematical terms, for instance if they are measurements of physical quantities. As this suggests, such practices are not confined to the *recherché* outer reaches of scientific practice and can obtain in quite mundane contexts. The ordinary macro-objects of everyday experience can be described as sets of points. To talk then of the 'nominalistic content' of theories is to talk of something very meagre indeed. If we take it that mathematicized science can tell us anything interesting about the world then we should hope that it is not only the nominalistic parts of mathematicized scientific theories that are confirmed. While Sober is surely right in thinking that, in actual scientific practice, the status of mathematical entities is not on a par with theoretical ones⁵, these structural features of mathematicized theories show that his argument ought not to be cashed out in the manner it is.

Sober treats M as though it is a purely mathematical theory, but, in fact, competing theories will contain different bridge laws, in which case one (partly) mathematical theory *could* be contrastively confirmed over another. Even holding bridge laws fixed, there are problems. Taking the case of measurement, 'The mass of d_1 is 5 kilogrammes' could be contrastively confirmed over 'The mass of d_1 is 5^{10} kilogrammes'. Both claims however entail the existence of mathematical objects. Because the mathematical part of a mathematicized theory cannot simply be held apart from its nominalistic part, the indispensability argument cannot be blocked by rejecting confirmation holism in the manner of Sober. At the same time, and for the same reasons, it is not clear that a great deal is gained for the advocate of the indispensability argument in *adopting* confirmation holism, *contra* Colyvan. If propositions such as 'The mass of d_1 is

⁵A case made painstakingly by Penelope Maddy [1997].

5 kilogrammes' are the kind of thing that is taken to be confirmed by observation, then controversial epistemological doctrines are not required for the claim that the existence of mathematical objects is confirmed by empirical means. To put things slightly differently, if propositions such as 'The mass of d_1 is 5 kilogrammes' are confirmed *directly*, then confirmation holism is not needed to get the platonistic argument off the ground.⁶ What about Leng? In fact, little of substance hangs on this point. For although Leng couches the debate in terms of confirmation holism, what emerges from her work is the thesis that although the non-nominalistic parts of our theories are, strictly speaking, false, they can be nominalistically adequate, and *this* is what is subject to empirical confirmation. This is a view I share, and one to which we will return shortly.

5.1.3 Premiss 1

So much for premiss 2. Premiss 1 has probably attracted more attention than any of the other premisses. Repudiating the indispensability thesis has been the motivation for the most prominent and celebrated nominalist research programmes, whose goal has been to systematically *expunge from our discourse* any mention of abstract objects, replacing it with talk of what *can* be constructed [Chihara 1990], what structures *might* have been [Hellman 1989], regions of space-time points [Field 1980], or even ideas in the mind of God (cf. Morris and Menzel [1986] for a relatively recent take on this and Plantinga [2012] for some contemporary support). Of these Field's programme has received by far the most discussion, and although Field himself has had little to say about it since the early Nineties, work continues to be done: Frank Arntzenius and Cian Dorr [2012] for instance have extended Field's nominalization from Newtonian mechanics to general relativity. How plausible is premiss 1? The claim is usually motivated by an appeal to the failure, so far, of philosophers to produce a strategy for nominalizing current science. As such, the argument takes the form:

Up until this point philosophers have been unable to nominalize our best scientific theories.

∴ Our best scientific theories cannot be nominalized.

But why think a thing like that? Why think that we can move from the *de facto* state of affairs, which regards the extent to which our theories have been nominalized, to the *modal* claim about nominalizability? What is required is something like a completeness claim with respect to our current state of knowledge on the subject, or perhaps a pessimistic induction from failed attempts to dispense with mathematical terms from physics. The former option is, I take it, roundly implausible, but what of the latter? Perhaps failure to nominalize physics so far confers some small degree of confirmation on the proposition that physics *cannot* be nominalized, but it is hard to see that these considerations should give one any great credence in the premiss.⁷ Consider an analogous example. There remains, at the time of writing, no

⁶Interestingly, Colyvan himself notes—at some occasions in the same breath as urging the importance of confirmation holism—that much of scientific theories are made up of mixed mathematical-physical sentences.

⁷Granted, there are specific problems with extending Field's own programme to phase space theories and quantum mechanics; but what is required to establish *indispensability* is that *no* such programme *could* be made to work.

proof of Goldbach's famous conjecture that every even integer greater than 2 can be written as the sum of two primes. To what extent should this observation confer confirmation on the claim that the conjecture is unprovable? It would be rash, to say the least, that we should be confident of the unprovability of the conjecture on those grounds; *recherché*, academic problems of this ilk are hugely difficult—an onerous task for even the most celebrated practitioners of the field—and there is little to recommend the thought that were a proof to be had, we would have likely discovered it. If anything the pessimistic induction required by the indispensability argument has even less to recommend it. There are accepted standards of proof for mathematical theorems, and no doubt more mathematicians have worked over more years on proving Goldbach's Conjecture than philosophers have on nominalizing physics. Even if one accepts the terms of the debate laid out by the platonist (i.e. that indispensability is central to the tenability of nominalism), the indispensability argument is inconclusive at best.

Whither Nominalization?

What though of the terms of the debate? What relationship does indispensability have to nominalism? The answer, so I say, is not a great deal. The indispensability of mathematical terms from our best theory of the world is neither necessary nor sufficient for a tenable nominalism. This goes contrary to the received wisdom amongst platonists, such as Colyvan, who holds that any viable form of nominalism 'require[s] the success of a hard-road strategy such as Hartry Field's nominalization program' [Colyvan 2010: 303]. Ramping up the rhetoric, if perhaps not the argumentation, Burgess puts it somewhat more memorably:

The main achievement of our book [*A Subject With No Object*] will have been to provide decent burial for the hard-working laborious variety of nominalism. For almost everything that has come forth since from the nominalist camp has represented the light-fingered larcenous variety, which helps itself to the utility of mathematics, while refusing to pay the price either of acknowledging that what mathematics appears to say is true, or of providing any reconstrual or reconstruction that would make it true. [Burgess 2004:18]

I take it here that Burgess thinks this is a bad thing, and whether or not one accepts a claim as strong as Colyvan's there is a felt sense in the literature that successful nominalization would be a better result for the nominalist than a strategy for explaining why mathematics finds such successful application in the sciences that does not appeal to nominalization. What suppositions are required to justify this view and should they compel our assent? To make explicit the reasoning involved in this claim we can begin by drawing a distinction between *ontological* nominalism and *logical* nominalism. The first is the thesis that abstract objects do not exist, the second is the thesis that our language can be expunged of reference to abstract objects without impairing our ability to describe the world (adding perhaps some other constraints on the nominalistic base language such as, for instance, that this expressive power is available to creatures with similar expressive and conceptual resources as *Homo sapiens*). The question of whether nominalists must nominalize then is really a question about the relationship between these

two distinct theses. Fundamentally, those committed to nominalization being required for nominalism hold to the following conditional:

(Ont→Log): If mathematical objects do not exist then it is possible to characterise the concrete domain without quantifying over them.

Arntzenious and Dorr are unusual in stating this sort of idea explicitly,⁸ but the import of doing so should not be underestimated; making the claim explicit serves to make it available to scrutiny. Really, (Ont→Log) is not what the advocate of nominalization is after, for the conditional may hold merely as a matter of contingency, and this would not tell us anything about whether or not there is some *relationship* between nominalism and nominalizability. What is needed is something stronger:

(□[Ont→Log]): Necessarily [If mathematical objects do not exist then it is possible to characterise the concrete domain without quantifying over them.]

Putting this picturesquely in the language of possible worlds, we would say that there are no possible worlds in which mathematical objects do not exist, but it is not possible to characterise the concrete domain without quantifying over them. But why think a thing like that? The advocate of (□[Ont→Log]) must maintain that there is a link between the existence of mathematical objects and our ability to describe the world without quantifying over them. Beginning with a possible scenario in which (i) mathematical objects exist and (ii) it is impossible to characterise the concrete domain without talking about them, (□[Ont→Log]) entails that were we to vary this scenario by making its mathematical domain empty, it would, *as a result*, become possible to characterise the concrete domain without quantifying over mathematical objects. But this is surely not the case, so (□[Ont→Log]) fails.

It may be objected here that to endorse this view is to sink into scepticism. If we frequently cannot say straightforwardly *true* things about the world, then we have no real grip on what the world is like. Yet, we do have a grip on what the world is like, so we must be able to say straightforwardly true things about the world. While it may be initially compelling it is, I think, ultimately without force. It is easy to construe the complaint as a rejection of a view we will call inarticulabilism:

Inarticulabilism: It is not possible to describe the world comprehensively, using straightforwardly true sentences.

The nominalist, it is thought, is committed to inarticulabilism, but inarticulabilism is unacceptable. The nominalist however, need make no such commitment. Only nominalism in conjunction with the thesis that our best scientific theories cannot be nominalized is committed to inarticulabilism, but, as we have seen, there are no compelling reasons to think that our best scientific theories cannot be nominalized. As far as the nominalist is concerned, whether inarticulabilism is true or not is something we could come to find out; we are free to follow the evidence where it leads. Nominalism *itself*—even that of the ‘light-fingered larcenous’ variety—involves no commitments about what we are or are not able to

⁸ ‘If [mathematical objects] don’t exist, shouldn’t it be possible, at least in principle, to characterise the physical world without talking about them at all?’ [Arntzenious and Dorr 2012: 3]

express. In fact, it is those who claim that we must be able to describe the world in straightforwardly true sentences that are committed to a thesis about what we are able to express:

Articulabilism: It is possible to describe the world comprehensively, using straightforwardly true sentences.

But how might we come to know such a thing? Is it something that we could come to know a priori? It is not a logical or conceptual truth, nor is it true by definition. In fact, the premise is clearly not a candidate for a priori knowledge. The premise is substantive, contingent, and, if it is possible at all to know whether something like articulabilism is true, it must be knowable a posteriori. Specifically, whether we can describe the world directly, using straightforwardly true sentences is a function of (i) the expressive powers of our language, (ii) the complexity of the world we wish to describe, and (iii) what exists. Finding out whether we can do so involves an assessment of (i)–(iii). With regards to (iii), if mathematical objects exist then we will, to a far greater extent, be able to describe the world using straightforwardly true sentences. Since however warrant for thinking that we can describe the world using straightforwardly true sentences depends on warrant for thinking that mathematical objects exist, this objection against the form of nominalism we are considering is circular. What the platonist needs then, if she is to attack the nominalist on these grounds, is something stronger:

Strong Articulabilism: Necessarily, it is possible to describe the world using straightforwardly true sentences.

Strong articulabilism however entails $(\Box\text{Ont} \rightarrow \text{Log})$ and so is subject to the same criticisms. The nominalist is not committed to inarticulabilism, but even if she were, there are no convincing reasons to think that we must be able to describe the world directly, with straightforwardly true sentences. Philosophers, being what they are, would like all facts to be articulable, but perhaps they are not. Moreover, the claim that if mathematical objects do not exist then we should be able to describe the world without reference to them, is untenable. Logical nominalism is not a necessary condition for ontological nominalism.⁹

Nor is it clear that nominalization is a *sufficient* condition for a tenable nominalism. This claim is characterized by the principle:

$(\text{Log} \rightarrow \text{Ont})$: If it is possible to characterise the concrete domain without quantifying over mathematical objects then mathematical objects do not exist.

Let Γ be a mathematicized theory and Γ^N its *nominalistic counterpart*. The question here is in what way our ability to produce Γ^N is relevant to the truth or falsehood of Γ . The platonist's claim is that Γ is confirmed by empirical evidence. Does the existence of Γ^N undermine this claim? It is not clear why this should be so. Perhaps the nominalist could claim that Γ^N is an *analysis* of Γ ; i.e. that it has the same *meaning* as Γ . However, there is a problem of *symmetry* here: if a mathematicized theory Γ and its nominalistic counterpart Γ^N mean the same thing, then what is to say that Γ^N provides the

⁹Later in the chapter I will say something about how it is possible to characterise the world in lieu of doing so with straightforwardly true sentences; *viz* why nominalism is not committed to a descriptive scepticism.

true meaning of Γ , rather than contrariwise? At any rate, that Γ and Γ^N mean the same thing is a difficult position to maintain. What one is confronted with is not akin to Russell's analysis of definite descriptions; where translating natural language locutions into predicate logic with identity served to clarify their truth-conditions. Γ and Γ^N are formalized theories whose truth-conditions they wear on their sleeves and which, on the face of it, appear to be different from one another: Γ is committed to the existence of mathematical objects and Γ^N is not. Nor can the same inferences be made with each. It may be the case that Γ is conservative *over* Γ^N , in that no inferences can be made in the vocabulary of Γ^N using $\Gamma + \Gamma^N$ that could not be made in Γ^N alone. But the two vocabularies are not inferentially equivalent *simpliciter*, since Γ licenses inferences about mathematical objects and Γ^N does not. It is true of course that neo-Fregeans *do* make a claim of this sort: the right and left hand side of the biconditionals in abstraction principles express "recarvings" of the same content. But neo-Fregeans have never been able to sufficiently demystify the appeal to recarvings to show how this might be the case, in the face of strong reasons to think otherwise. Perhaps there is some other, weaker equivalence relation \sim that holds between Γ and Γ^N , which would justify taking one to express the ontological commitments of the other. But then the symmetry problem arises again, for equivalence relations are symmetrical. Another kind of option is to claim that the nominalistic theory Γ^N is ultimately *superior* to Γ —perhaps on the grounds that it is more ontologically parsimonious—and hence that one should reject Γ in favour of Γ^N . This too is problematic; for Γ and Γ^N are not *competing* theories, they are not *incompatible* with one another. As such, confirmation of one theory does not constitute *disconfirmation* of the other. Another route is to note—as Field does (indeed *proves*) in the case of Newtonian gravitational theory—that Γ is a conservative extension of Γ^N . Since Γ^N is true this might be taken to explain the predictive success of the (according to the nominalist) false Γ . But is there such an association between our expressive powers and ontological questions? Suppose another theory Γ^S was created, the language of which referred to and quantified over only sense data, and that it was then shown that Γ^N was a conservative extension of Γ^S —i.e. that all the consequences of Γ^N expressible in the language of Γ^S are consequences of Γ^S . Would this be sufficient to show that concrete objects did not exist, or that we are not warranted in taking them to exist? We do not believe in concreta because we indispensably quantify over them, nor should we; there are other reasons for taking concrete objects to exist. Nominalization is not necessary for a defence of nominalism, and nor does it appear to be sufficient.

Heavy-Duty Platonism

We have described the nominalist's situation as this: when she cannot make true utterances about the world, she must resort to making merely nominalistically adequate utterances, and by doing so *gesture* indirectly at what the world is like. This predicament however is not peculiar to the nominalist. Consider one representative application of mathematics: measurement. Platonists hold that there are relations between concrete objects and numbers; for instance, there are an infinity of functions M which represent the mass of any object a as a real number M_a . A single concrete object which bears the *mass-in-*

kilogrammes relation to the real number 100 will also bear the *mass-in-stones* relation to the real number 15.75. On the usual view these relational facts are not fundamental; they hold in virtue of fundamental facts about concrete objects and numbers: our choice of measurement scale is arbitrary, but once we have chosen a scale the real number which represents the mass of a is determined by the intrinsic properties of a . Were the intrinsic properties of a different, the real number which represents its mass would also be different. This picture of measurement presupposes there being purely concrete objects and their having purely concrete properties. What it doesn't entail is that we can provide intrinsic descriptions of these objects; if they are complex enough it may be easier (or necessary) to represent these intrinsic properties by assigning them numerical values. While this platonistic view admits of our mixed sentences being true, it is in precisely the same boat as the view presently under consideration when it comes to our grip on what the world is like. Even if we countenance the existence of abstract objects, we remain committed to our inability to directly articulate facts about the concrete domain. As in the present view, we gesture at what the concrete domain is like by describing relations that hold between it and the abstract domain. Platonist and nominalist alike, only those who think we can nominalize our discourse can consistently hold that we can directly express facts about physical objects and systems.

Both ordinary nominalists and ordinary platonists are committed to the view that, without nominalizing our best theories, we are unable to directly describe the purely concrete facts. But there is another kind of view which can avoid this. *Heavy duty platonism* takes a different line on the relations between concrete and abstract objects, as Field explains:

According to moderate platonism, *such magnitude relations between physical things and numbers are conventional relations that are derivative from more basic relations that hold among physical things alone*. The heavy duty platonist rejects this, taking the relation between physical things and numbers to be a brute fact, not explainable in other terms. (If one likes to flaunt one's heavy duty metaphysics, one can say that there is a mysterious relation of platonic participation between physical things and numbers. But the position is the same whether or not one flaunts it.) [Field 1989: 186]

For the heavy duty platonist what is ontologically fundamental are not concreta and the relations that hold between them (indeed no such relations exist) but mathematico-physical relations. For the heavy duty platonist there are no concrete facts, only mixed facts. The heavy duty platonist *can* avoid the charge of mysticism: there is no problem in grasping the properties of the concrete domain that lie behind our mixed sentences because *there are no* concrete properties; the world is a mathematical-physical alloy. The heavy-duty platonist will see no value in nominalization: if the concrete world does not have an *intrinsic* structure, but only relations to the mathematical domain, then a nominalization of physics—providing, as it does, a description of the intrinsic structure of the concrete domain (or an intrinsic description of the structure of the concrete domain)—serves no purpose.

The heavy-duty platonist may motivate their view on the grounds that it is the only view that does justice to the logical structure of scientific theories. As we have just noted, many theories will describe

the world in terms of relations between concrete and mathematical objects, *not* in terms of the intrinsic structure of concrete systems.¹⁰ The scientific picture of reality, so says the heavy-duty platonist, *is* of a world which consists of inherently mixed facts; physical-mathematical relations. To suggest that there is a concrete domain with interesting properties and relations underlying this is to reject the picture of the world that contemporary science gives us. Despite this *prima facie* support, heavy-duty platonism should be resisted. Mark Balaguer discusses one reason why this is so, taking as an example the apparently mixed fact:

(A) The physical system S is forty degrees Celsius.

He states:

The person who objects in this way fails to appreciate the full significance of the causal inertness of mathematical objects. It is no doubt true that (A) says that S stands in the Celsius relation to the number 40. But since 40 isn't causally relevant to S's temperature, it follows that if (A) is true, it is true in virtue of facts about S and 40 that are entirely independent of one another, that is, that hold or don't hold independently of one another. In other words, if we grant that the number 40 isn't causally related to S—and this is beyond doubt—then we are forced to say that while (A) does express a mixed fact, it does not express a bottom-level mixed fact; that is, the mixed fact that (A) expresses supervenes on more basic facts that are not mixed. In particular, it supervenes on a purely physical fact about S and a purely platonistic fact about the number 40. [Balaguer 1998: 133]

The thought here is that because temperatures are causally relevant—a system with a temperature of 40°C has different counterfactual properties to a similar system with a temperature of 30°C—but numbers are not, there must be purely concrete facts underlying this mixed fact that account for the causal or counterfactual differences between a 30°C system and a 40°C system. Balaguer concludes that 'some purely physical fact that involves S holding up its end of the "(A) bargain" obtains' [ibid., pp.133-4]. Although Balaguer focusses on the acausality of abstract objects, this is but one argument in a taxon of similar considerations:

Modal Objection: Mathematical objects, if they exist, have their properties necessarily, but mixed facts are contingent. As such these mixed facts must be grounded in contingent, purely concrete facts.

Temporal/Spatial Objection: Mathematical objects, if they exist, do not have temporal or spatial locations, but mixed facts do. As such these mixed facts must be grounded in temporally and spatially located, purely concrete facts.

Event Objection: There are no mathematical events, but there are mixed events. As such these mixed events must be grounded purely concrete events.

Change Objection: Mathematical objects, if they exist, are changeless, but mixed facts change over time. As such, this change must be grounded in purely concrete facts.

¹⁰Set aside for the time being the complication for the heavy-duty platonist that a significant portion of good contemporary science does not describe the world mathematically.

Epistemological Objection: Mathematical facts are knowable a priori, but mixed facts are knowable a posteriori. As such, what is discovered a posteriori about mixed facts must pertain to purely concrete a posteriori-knowable facts.

These arguments may not speak decisively against heavy-duty platonism. At best they provide a *rebutting* defeater; they give us reason to think that heavy-duty platonism is false. They do not, however, address directly the original considerations about the logical structure of scientific theories, thought to give countenance to the doctrine. An *undercutting* defeater¹¹ for heavy-duty platonism would give us reason to doubt the cogency of *this*, and in the absence of an undercutting defeater the two sides may remain at loggerheads.¹² A closer look at the structure of mathematized theories tells us that, *pace* the heavy-duty platonist, such theories are committed to concrete systems having intrinsic concrete properties. We can see this by noting that measurements are expressed in terms of *functions* from concrete systems to real numbers. The domain of a function F is the set $\{x : \text{there is some } y \text{ such that } \langle x, y \rangle \in F\}$, and the range of a function F is the set $\{y : \text{there is some } x \text{ such that } \langle x, y \rangle \in F\}$. What this means intuitively is that the function takes elements of the domain as inputs; when the function is one from concrete systems to real numbers, as in temperature ascriptions, the domain is the set of relevant concrete systems and the range \mathbb{R} ; when the function is from space-time points to vectors, as in descriptions of magnetic fields, then the domain is the set of space-time points and the range the set of vectors. The value of the function is determined by what it takes as its argument, and, in the case at hand this means that the *properties* of the space-time point determine the value of the function; so there are concrete entities with concrete properties. More pertinent to the point at hand, it cannot be that the object taken as the argument determines the value of the function *regardless of its own properties*; as would be entailed by a view such as heavy-duty platonism. For there is nothing about being the space-time point designated by, for instance, $\langle 7, 9, 12, 4 \rangle$ that would result in the function from that space-time point giving one value rather than another. Rather, it must be the properties of that space-time point that result in the function outputting the value it does. We are left with a picture of the world where fundamental concrete objects have fundamental concrete properties which are *designated* by numerical magnitudes.

5.1.4 Premiss 3

So much for heavy-duty platonism, and for premisses 1, 2 and 4 of the indispensability argument. Quine's criterion is not here under dispute. As for confirmation holism and the indispensability thesis, I have argued that the former is something of a red herring, and that the latter not only lacks support, but is strictly orthogonal to a defence of nominalism. What of the third premiss?

3. *Theory Realism / Naturalism:* We look to our best scientific theories as the ultimate arbiter of

¹¹Cf. Pollock [1986], Bergmann [2005] and Plantinga [unpublished] for a discussion of the difference between rebutting and undercutting defeaters.

¹²The entire nominalism/platonism debate could well be seen in such terms. Platonists may be troubled by epistemological objections to their view, but so long as they take indispensability arguments to be cogent they will, at best, see the situation as a stand-off between the two doctrines.

existence and truth.

The premiss obscures more than it illuminates. How are our best scientific theories to be used to determine what exists? What does this involve?

Quine, Realism and Anti-Realism

Quine provides one answer to this question. The Quinian recipe for ontological commitment is striking for its simplicity. We begin by taking our best theory of the world, and translate it into a canonical formal language—one whose domain of quantification is wholly unambiguous. Thus equipped, we simply read what exists off the face of the theory; what exists is what the theory says exists, or, equivalently, what is required for the theory to be true.¹³ No surprise then that the method has found such wide acceptance; it appears close to tautological: what we should take to exist is what our best scientific theory of the world says exists.¹⁴ Not only is the method *prima facie* plausible; it also promises to take ontology out of the hands of metaphysicians—thus liberating it from internecine debates about metaphysical intuitions—and places it into the hands of the scientific community. Ontology, on this picture, can be settled empirically, and progresses as scientific knowledge progresses. This picture of ontology has, to a large degree, stuck, and ossified into the assumed substructure of many subsequent debates. In the philosophy of mathematics, for instance, it has been the motivation for the most prominent nominalist research programmes, whose goal, as we have already discussed, has been to expunge from our discourse any mention of abstract objects, replacing it with talk of what *can* be constructed, what structures *might* have been, regions of space-time points, or ideas in the mind of God.

Yet despite its widespread acceptance and pervasive influence over the philosophical landscape, the Quinian method of ontology is afflicted with at least two congenital problems. The first is that it is inherently *anti-realist*. (The second is that it mislocates the locus of ontological commitment; but this is an issue that we will address later in the chapter.) There is some irony in the evocation of Quine in order to settle the *first-order* ontological dispute over the existence of abstract objects, for Quine himself was not an ontological realist. Quine's view is usually contrasted with that of his mentor Carnap who held that there were meaningful existence claims only within a linguistic framework—which could be determined either empirically or analytically—and that the choice of which framework to accept was a pragmatic one; but Quine himself shared this metaphysical pragmatism, differing from Carnap in this regard only insofar as he took it that questions could not be partitioned into those internal to a linguistic framework and those external to a linguistic framework; since the demarcation of separate linguistic frameworks relied on the analytic-synthetic distinction, a distinction that Quine gainsaid. Thus for Quine, the practice of making pragmatic decisions about which ontology to adopt was not limited to choosing convenient linguistic frameworks but was diffused throughout our entire web of belief. Quine

¹³This is a slight simplification. Quine held also that this *prima facie* ontological commitment was *defeasible* in those cases where the quantification at hand was ultimately *dispensable*. Indeed, this qualification drives the indispensability argument itself.

¹⁴One stripe of naturalist will add *and only what our best scientific theory of the world says exists*.

was a Carnapian.¹⁵

The problem for the Quinian method of ontology in the context of first-order debates about the existence of various classes of entities, is that some kind of ontological anti-realism seems difficult to avoid, indeed *built-in* to the method itself. A theory, as logicians understand the term, is a set of sentences closed under entailment; but there are many different ways to make a formula or set of formulae true, many different models of a given theory. To illustrate, consider a theory Γ which states $\exists x \exists y \exists z (Rxy \wedge Ryz \wedge Ryz)$, and a structure $\mathcal{S} = [D = \{a, b, c\}, R = \{\langle a, b \rangle, \langle b, c \rangle, \langle a, c \rangle\}]$ with domain D and relations R . Any world with this structure satisfies Γ . This could be a world in which a, b and c are rocks, and R is the *larger than* relation, or a world in which a, b and c are cheeses, and R the *more pungent than* relation. Any number of worlds, with entirely different ontologies, exemplify \mathcal{S} and satisfy Γ . In this sense truth comes first, and ontology is secondary, for there are many different ways of analysing the sentences of a theory, such that their truth values will remain intact. Quine even provides a recipe for doing so: a one-to-one function f is defined over the objects of the relevant domain \mathcal{D} that maps every object o onto its spatiotemporal complement (all of space-time excluding o). Every sentence that appears to refer to o is reinterpreted as referring to $f(o)$ and every predicate is reinterpreted so that it holds of $f(o)$ if and only if it previously held of o . Not only are there multiple models of a given consistent theory, given the downward Löwenheim-Skolem theorem which says that if a theory has infinite models then it has models that are *countable*, all that is required to secure the truth of our best theory of the world is that there be denumerably many things, be they sets, integers or wheels of liverot. In answer to the question *What does our best science say exists?*, one can reply *Infinitely many wheels of liverot*. Jonathan Schaffer, in a discussion of this point, puts it well when he says ‘Such a view invites the reply: if that was the answer, what was supposed to be the question?’ [Schaffer 2009: 350].

It may be hoped by the Quinian that she can circumvent this unhappy consequence by claiming that the natural language locutions that have been regimented into predicate logic, are in fact what fixes the referents of our best theory of the world. Doing so would involve forfeiting some of the benefits that regimentation was intended to supply—formalizing our best theories was supposed to *clarify* their ontological commitments. It would also require denying what Quine called the ‘inscrutability of reference’.¹⁶ Quine held that all facts were physical facts, and hence that any facts not specified in the language of physics must themselves be *determined* by physical facts. Yet, facts about reference are not expressed in the language of physics, nor determined by the set of all physical facts, so there are no determinate facts about reference [Quine 1960]. However, either one of these premises could be resisted. In his writings, Quine fails to clarify the determination relation involved in these claims. Scott Soames [2003] makes a forceful case that, depending on how we precisify the determination relation,

¹⁵This point is made by Huw Price [1997, 2009]: ‘Quine is not returning to the kind of metaphysics rejected by the logical empiricists. On the contrary, he is moving forwards, embracing a more thoroughgoing post-positivistic pragmatism. In this respect, far from blocking Carnap’s drive towards a more pragmatic, less metaphysical destination, Quine simply overtakes him, and pushes further in the same direction.’ [Price 2009: 327].

¹⁶Mentioning, as opposed to using, this epithet is apt, as it is something of a misnomer; the doctrine being *metaphysical* rather than epistemic in character.

physicalism may be plausible, and the underdetermination of reference may be plausible, but that there is no precisification in which both are plausible.¹⁷ If the determination relation is taken to consist in logical consequence then certainly facts about reference are not a logical consequence of any theory found in physics, indeed trivially so, as the vocabulary of theories of reference contains predicates and terms not found in the vocabulary of physics; the underdetermination thesis holds. However, the truths for instance of biology, or for that matter the everyday deliverances of our experience such as *that house has a thatched roof*, also include vocabulary not found in physics and so are also not logical consequences of physics. It is implausible though that the claims of biology, or the deliverances of experience are indeterminate. Nor will bolstering the determination relation to one of *a priori* consequence by ameliorating it with *a priori* definitions and principles help, as the facts of biology and everyday deliverances of experience would still not be determined by the facts of physics. This is because for a theory to be an *a priori* consequence of physics, *a priori* bridge principles between the vocabulary of physics and the theory must be established, but the kind of bridge principles that relate results from different areas of science are established *a posteriori*.

These precisifications are broadly epistemic in character, but one can parse the determination relation differently, as a metaphysical relation between facts. We might say that a set of statements Γ determines another set of statements Δ if and only if it is necessary that if all the statements in Γ are true then all the statements in Δ are true. On this interpretation physicalism is made more plausible—it is impossible, we might think, that two worlds could differ with respect to the biological facts without differing with respect to the physical facts—but at the cost of making the underdetermination of reference less plausible (indeed for the same kinds of reasons that, given this determination relation, physicalism is made *more* plausible). For facts of reference to be underdetermined by physical facts in this metaphysical sense, one must suppose that two physical duplicates, with identical brain states, in physically identical worlds, with physically identical relations to the linguistic community and environment, could refer to different things on making identical utterances of the word “rabbit”. This is implausible, or at any rate difficult to assess with any degree of certitude. Let us assume then for the time being that the proponent of the Quinian method of ontology need not thereby be committed to the underdetermination of reference, and, as a consequence, the doctrine of ontological relativity.

Even given this (plausible) assumption, the Quinian method cannot easily avoid anti-realism, for there are two roads to ontological anti-realism from the Quinian method of ontology. The first of these stems from the plurality of models of any given theory, but the second arises from an inherent structural feature of mathematicized theories. Our best theory of the world is a somewhat malleable thing; a point that Quine readily recognised in a famous passage:

What then is the brave new ontology? There are the real numbers, needed to measure the intensity of the various states, and there are the space-time regions to which the states are

¹⁷Strictly, Soames is discussing the underdetermination of translation, but, given that the underdetermination of reference is a corollary of this, the same considerations apply.

ascribed. By identifying each space-time point with a quadruple of real or complex numbers according to an arbitrary system of coordinates, we can explain the space-time regions as sets of quadruples of numbers. The numbers themselves can be constructed within set theory in known ways, and indeed in pure set theory; that is, set theory with no individuals as ground elements, set theory devoid of concrete objects. The brave new ontology is, in short, the purely abstract ontology of pure set theory, pure mathematics. [Quine 1978: 164]

There are many different ways of expressing our best theory of the world; ways which quantify over wholly different domains. Given Quinianism, there is an intrinsic and ineliminable element of choice in what exists. Of course this is no problem for a metaphysical anti-realist. Quine himself was well aware of this consequence, opining ‘ontology is not what mainly matters’ [ibid.], or, more vividly:

Reference and ontology recede thus to the status of mere auxiliaries. True sentences, observational and theoretical, are the alpha and omega of the scientific enterprise. They are related by structure, and objects figure as mere nodes of the structure. What particular objects there may be is indifferent to the truth of observation sentences, indifferent to the support they lend to the theoretical sentences, indifferent to the success of the theory in its predictions. [Quine 1990: 31]

Most philosophers however *are not* metaphysical anti-realists, yet endorse, either tacitly or explicitly, Quine’s ontological method, and therein lies a problem. Quine’s conclusions, whatever we make of them, were consistent with his method of ontology—he embraced (rather sanguinely, one might think) the drastically deflationary consequences of his doctrine—but the superficial Quinian pays lip-service to this process of determining what exists, whilst simply ignoring its radical anti-realist consequences.

Truth and Realism

‘Realism’ of course is something of a fraught term. The word can be used to locate views on at least two distinct axes of commitment. One—the more metaphysical of the realisms—concerns the independence of the world from the minds that inhabit it, and, relatedly, whether the language we use to describe the domain in question acts to pick out objects with particular properties in the manner suggested by the structure of our language. Dummett has this kind of concern in mind when he characterizes realism in the following way:

Each of the disputes concerns the interpretation of a certain class of statements; the realist maintains that these statements are to be taken at face value, while his opponent argues that what makes them true or false is not what the way they are framed would suggest, but something different. [...] Do statements about mental events and states describe conditions within worlds private to each one of us, or are they oblique ways of talking about our behaviour? Do mathematical statements depict how things stand in a special, nonphysical, unchanging sector of reality, or are they about the constructions that mathematicians can carry out? Are statements about the external, physical world made true or false by the

disposition of material objects existing independently of us and of our knowledge, or are they based entirely upon our sense impressions?

The answers to the metaphysical questions about this or that variety of realism thus turn on the correct interpretation of one or another class of statements. Metaphysics accordingly rests on semantics.' [Dummett 2010: 125]

Intuitionism and *game formalism* are both anti-realist in this sense. The other more epistemological axis concerns our ability to get things right about the way the world is. Fictionalists are anti-realists in this sense, regarding mathematical discourse. Even regarding this kind of realism, the fictionalist may be an anti-realist only in a rather attenuated sense; for while she may hold that there is no independently existing abstract domain that our mathematical discourse accurately describes, she need not hold that mathematicians are in *error*, or embroiled in an epistemic catastrophe, for she may not hold that the *goal* of mathematics is to describe an independently existing abstract domain. The *telos* of an activity may not always be flatly read off its semantics. Nevertheless, that mathematical statements are, strictly speaking, false, and scientific statements (partly) mathematical, is often thought to have the anti-realist consequence that mathematicized science cannot tell us what the world is like:

[C]onsider the argument that claims that ordinary mathematical and scientific judgements are problematic because they could only be true by a lucky accident. Surely, if ordinary views genuinely are problematic for this reason, they must be so whether or not there are available any alternative views one might adopt to replace them. [...] It would seem that if one reposes any serious confidence in the arguments in the negative, destructive side of the nominalistic literature, one ought to draw the conclusion that standard science and mathematics *are no reliable guides to what there is*. [Burgess and Rosen 1997: 60-61, my emphasis]

We should dwell on this for a moment, for this kind of consideration, I believe, has been critical in conditioning the nominalist-platonist debate. It is the same thought found in Lewis' credo, where he berates the nominalist for propounding a view of comical immodesty; the thought that nominalist entails that 'standard science and mathematics are no reliable guides to what there is'. This then is what we are faced with: either accept nominalism and reject the truth of our best scientific theories (and hence the claim of good science to be a guide to what there is), or retain the truth of our best scientific theories (and hence the claim of good science to be a guide to what there is) and reject nominalism. No wonder then, given this stark dichotomy between nominalism and science as a guide to reality, that most philosophers are unwilling to embrace nominalism. The nominalist becomes rather like 'the sceptic' of epistemological discourse; a spectral figure haunting philosophical tracts, and providing a handy foil to more acceptable views, but rarely manifesting herself as a flesh and blood philosopher. The nominalist can level any argument she chooses, but if the conclusion is that science is no reliable guide to what there is then her *modus ponens* will always be, to the philosopher of a more moderate disposition, a *modus tollens*. For that matter, if I held to this dichotomy, I would also disclaim my pretheoretical nominalism, rather as one, in becoming aware of quantum mechanics, learns to eschew,

perhaps reluctantly, determinism or intuitive notions of location. The dichotomy however, is a false one. Platonism, to be sure, secures us scientific realism in the sense that our best scientific theories may be true, but as we have seen, ‘realism’ of this sort is well-nigh powerless to tell us anything of interest about the nature of the world. Posing the question ‘Do you hold that our best scientific theories are true?’, the realist can reply ‘Yes, for I hold that there are denumerably many wheels of livarot, and nothing else’. I take it that merely knowing the *cardinality* of reality does not amount to a ‘reliable guide to what there is’ in the sense that we would plausibly attribute to good science. ‘Realism’ of the sort secured by platonism, and excluded by nominalism, is not sufficient, to borrow Quine’s elegant phrase, for ‘limning the true and ultimate structure of reality’ [Quine 1960: 220]. But is realism of this sort a *necessary* condition for this task? I say it is not, but answering this question will require some discussion of how mathematics is applied; a topic to which we shall turn shortly.

Scientific realism of the Quinian sort is compatible with metaphysical anti-realism of the Quinian sort, and hence has little to do with according good science its proper status as a source of knowledge and understanding of the natural world. Moreover, since the Quinian method of ontology entails an ontological relativism, it is not appropriate to invoke it in first-order ontological debates set against a backdrop of ontological realism. The issue of realism has been obscured by being couched in terms of truth. Truth was supposed to matter because we understood our ability to accurately describe the world in terms of truth, and because the platonist had it and the nominalist did not. In fact, the truth of our scientific theories turned out to have little to do with our ability to accurately describe the world. To gain traction on the issue requires a more generous methodology; one that factors in more facets of the scientific enterprise than the truth or falsehood of the sentences that compose our best scientific theories. *Pragmatic* features of science must be taken into account.

5.2 Indispensability and Pragmatism

Putnam is credited along with Quine as being an early proponent of indispensability arguments for the existence of mathematical objects. For instance, in his paper *Philosophy of Logic* he makes an argument centering around indispensability:

So far I have been developing an argument for realism along roughly the following lines: quantification over mathematical entities is indispensable for science, both formal and physical; therefore we should accept such quantification; but this commits us to accepting the existence of the mathematical entities in question. This type of argument stems, of course, from Quine, who has for years stressed both the indispensability of quantification over mathematical entities and the intellectual dishonesty of denying the existence of what one daily presupposes. [Putnam 1971: 347]

Above, I have tried to show that the link between the *indispensability* of mathematical terms in our scientific discourse and our warrant for taking them to exist is, at best, a loose one. Putnam however

made a different kind of argument for the existence of mathematical objects, one that shifts, in one way, from the *semantic* issue of the *truth-status* of mathematical sentences in our best scientific theories, to the *pragmatic* issue of the *role* that mathematics plays in our theorizing. The nominalist is committed to a certain kind of instrumentalism, insofar as she rejects the truth of mathematicized theories,¹⁸ but instrumentalism of this sort may be problematic for reasons orthogonal to indispensability. The thought here is that instrumentalism makes a miracle of the success of science. Jack Smart makes a classic statement of this idea:

If the phenomenalist about theoretical entities is correct, we must believe in a *cosmic coincidence*. That is, if this is so, statements about electrons, etc., are of only instrumental value: they simply enable us to predict phenomena on the level of galvanometers and cloud chambers. They do nothing to remove the surprising character of these phenomena. [...] Is it not odd that the phenomena of the world should be such as to make a purely instrumental theory true? On the other hand, if we interpret a theory in a realist way, then we have no need for such a cosmic coincidence: it is not surprising that galvanometers and cloud chambers behave in the sort of way they do, for if there really are electrons etc., this is just what we should expect. A lot of surprising facts no longer seem surprising. [Smart 1968: 39]

Here Smart is addressing instrumentalists about theoretical or unobservable entities, but the issue has a bearing on our discussion too. Putnam is aware of this and suggests a sister argument against mathematical nominalism:

I believe that the positive argument for realism has an analogue in the case of mathematical realism. Here too, I believe, realism is the only philosophy that doesn't make the success of science a miracle. [Putnam, 1971: 73]

Arguments of this sort, as I have been emphasizing, are very different in character from indispensability arguments. Indispensability arguments focus on the (indispensable) *presence* of quantification over mathematical objects in our best theories of the world, Putnam's no-miracle argument shifts the terms of the debate onto the way in which we *use* theories and what it is possible to achieve by using theories in this way. This, I think, is a far more promising avenue for the platonist to pursue. The Quinian method is based on an incomplete understanding of our 'best theory of the world'. While it attends to the semantic and syntactic aspects of theories, *pragmatics* is left unheeded. This is crucial because, as I will argue below, how a theory is *used* is essential to determining our ontological commitments. While a sentence or theory may be said to have ontological commitments *itself*, in some slightly extended sense, what is of interest to the ontologist is not the ontological commitments of a theory, but rather the ontological commitments one accrues by using a theory. In order to determine our ontological commitments we must locate the theories we use in the constellation of practices in which we use them. Putnam's no-miracle argument is an improvement over his indispensability argument precisely because it begins to do exactly this.

¹⁸Azzouni and Priest are exceptions, but recall that we rejected their non-ontologically committal account of quantification.

There are here two issues that ought to be kept distinct. One is the argument nominalism makes ‘the success of science a miracle’, and the other is whether nominalists can explain how mathematics comes to be successfully applied. Achieving the latter would have the effect of providing a rejoinder to Putnam’s argument, and we shall address this second issue when we examine the application of mathematics. However, unpacking Putnam’s rather terse formulation will allow us to state a quick, but less explanatory, riposte to the argument. For any theory x , let Sx be ‘ x is successful’ and Tx be ‘ x is true’. Let $\neg A$ be the negation of A , and $\Pr(A|B)$ be the probability of A given B . The argument runs thus: Theory h is very likely successful. If h were true, there would be a high probability that h was successful. If h were false there would be a low probability that h would be successful. Therefore there is a high probability that h is true. This can be stated formally as:

- (1) $\Pr(Sh) \gg 0$
- (2) $\Pr(Sh|Th) \gg 0$
- (3) $\Pr(\neg Sh|Th) \ll 1$
- (4) $\therefore \Pr(Th|Sh) \gg 0$

The argument enjoys a certain intuitive plausibility. Truth, however, is not a necessary condition for the success of a theory. As James Robert Brown notes, [1998: 1148] Ptolemaic astronomy predicted eclipses with an impressive degree of accuracy, and aether theory correctly predicted the surprising result that a shadow cast by a disk would have a light spot in its centre. Nor is truth a sufficient condition for the success of a theory, should false auxiliary assumptions combine with the (true) theory to produce false predictions. However, true theories (with true auxiliary hypotheses) do not make false predictions and only a very small proportion of false theories will make true predictions, and, on this basis, we infer that successful theories are more likely to be true than false. This conclusion is too hasty however and is obtained only if we ignore the plenitude of false and successful theories, relative to the comparatively meagre number of true successful theories. That is, because there are so many false theories relative to true theories, given any successful theory, that theory is more likely to be false than to be true. We draw the conclusion that successful theories are more likely to be true by being duped into making a base rate fallacy (cf. Howson [2000: 54]). As Peter Lipton explains:

[W]hen Putnam memorably says that it would take a miracle for a false theory to make all those true predictions, this way of putting it directs our attention towards the very low proportion of successful theories among false theories, but this has the effect of diverting our attention away from the low proportion of true theories among successful theories that is more to the point. [Lipton 2002: 197]

Letting \mathcal{H} be the base set of theories we can now reformulate the argument for all x , drawing out this hidden factor:

$$(5) \Pr(S|x \in \mathcal{H}) \gg 0$$

$$(6) \Pr(S|x \in \mathcal{H}) \gg 0$$

$$(7) \Pr(\neg S|x \in \mathcal{H}) \ll 1$$

$$(8) \therefore \Pr(T|x \in \mathcal{H}) \gg 0$$

The point here is that if, as Lipton asserts, $\Pr(T|x \in \mathcal{H})$ is low then we cannot conclude that a given successful theory is likely to be true. But is our base set really so full of false, but successful theories? Perhaps if every empirically adequate theory is included in \mathcal{H} then the answer is yes, but not every empirically adequate theory is such that we might be inclined to accept it. Empirical adequacy is not, after all, the only yardstick by which a theory is measured; we appeal to beauty, symmetry, simplicity, unificatory power and other such factors when assessing a theory. Perhaps most false successful theories are so *ad hoc* and unlovely that we naturally filter them out, and perhaps these filtering processes are themselves, somehow, truth-conducive. Perhaps then, those theories which pass through this gauze are so few in number, the base rate of successful false theories so diminished that, were a theory successful, we would have good reason to believe that it is true. Supposing that these caveats can be defended then Smart's miracle argument may be cogent in the case of theoretical entities. But what are we to make of Putnam's assertion that the argument translates to the case of mathematical objects? Notice that the miracle argument relies on the likelihood of a theory's truth being greater if that theory is successful (i.e. empirically adequate, capable of making novel predictions etc.). But we have already seen that mathematical objects do not make a difference to the physical world, which is to say that the truth about mathematical objects, whatever that may be, has no bearing on our empirical data. In which case we should endorse the following principle, where M is the hypothesis that mathematical objects exist:

$$\Pr(S|M) = \Pr(S|\neg M)$$

Informally, this means that so long as a theory is false *only in virtue of mathematical objects failing to exist* then its success is no less probable than if the theory were true. In other words, merely nominalistically adequate theories are no less likely to be empirically adequate than straightforwardly true theories. Hence, we must reject Putnam's claim that the miracle argument for realism about theoretical entities has an analogue in the case of mathematical realism¹⁹. Putnam's argument has been echoed elsewhere in the literature. Burgess, for instance, criticizes non-reconstructive nominalism on similar grounds, describing it as:

¹⁹In fact, Magnus and Callender [2004] reject the cogency of such 'wholesale' arguments for scientific realism, insisting that only 'retail' arguments can be used to determine the status of specific scientific conjectures, citing the debate over the existence of atoms which took place at the turn of the twentieth century as an example of such a retail argument. Hacking [1982] holds a related view. Note that these retail arguments rely on specific experiments and inference to the most probable cause (cf. Cartwright [1983], and Maddy [1997]), techniques which, because of the abstractness of mathematical objects, could not be used to determine their existence.

the unambitious kind that contents itself with saying that everything in the concrete world happens as if orthodox scientific theories mentioning mathematical objects were true, rather than the energetic kind that attempts to reconstruct those theories in a way that avoids mathematical objects. [Burgess 2001: 81]

When Burgess states that ‘everything in the concrete world happens as if orthodox scientific theories mentioning mathematical objects were true’, this is equivalent to the claim that everything in the world happens as if mathematical objects *exist*. Stated this baldly however the claim can be seen to have no force, for nothing in the concrete world happens *as if* mathematical objects exist; or rather, everything in the concrete world happens as if mathematical objects exist and, concurrently, as if they do not. This is what it is to say that mathematical objects are *abstract*. *Truth* cannot be that which makes the difference between a successful and unsuccessful theory, or that which explains the success or failure of a theory. The story to be told about this is more nuanced. Once we have explored the application of mathematics in greater detail we will be in a position to see more clearly not only *that* there is no conflict between nominalism and the empirical success of mathematicized theories, but also *why* the nominalist can adequately account for the empirical success of mathematics.

5.2.1 The Enhanced Indispensability Argument

So far I have been unspecific regarding the place of pragmatism in a defence of nominalism, but the role of practice in determining ontological commitments can be illustrated in the light of a recent debate in the philosophy of mathematics. The controversy centres around the theoretical benefits afforded by mathematics to empirical science, and whether they can be said to be the same *kind* of theoretical benefits that are gained through the postulation of concreta. If the conferment of a certain theoretical virtue through the postulation of a concrete object is grounds enough to include that object in our ontology then, if we are not to be accused of metaphysical chauvinism, the conferment of a certain theoretical virtue through the postulation of an abstract object is grounds to include *that* object in our ontology. Melia [2000] argues that the theoretical benefits gained from mathematics are not of the same sort as those gained from the postulation of concreta, and for this reason we are not compelled to accept the existence of those mathematical objects over which we quantify. For instance, simplicity is generally regarded as a theoretical virtue. Melia contends that, although mathematics may allow for the formulation of more simple theories about the world, it has not been used to demonstrate that the world itself is a simpler place. Theoretical entities, on the other hand, are used to offer unified and simple accounts of how the world is, such that it produces the phenomena that we seek to explain. Unification is another theoretical virtue, and Colyvan [2002] provides an example of complex analysis being used to unify physical theories which employ differential equations. Melia [2002] responds, arguing that Colyvan’s example does not represent mathematics being used to exhibit unification in the world, as ‘having a unified approach to solving equations seems quite a different matter from having a unified account

of apparently disparate phenomena' [ibid.: 77].²⁰ Baker [2005, 2009] intercedes, taking up Melia's challenge to find an example from science where 'the postulation of mathematical objects results in an increase in the same kind of utility as that provided by the postulation of theoretical entities' [Melia 2002: 75]. The utility in question is *explanatory utility*: given a mathematical explanation of some phenomenon, we are warranted in postulating the existence of the mathematical objects that feature in the explanation. Baker suggests an 'Enhanced Indispensability Argument':²¹

- (P1) We ought rationally to believe in the existence of any entity that plays an indispensable explanatory role in our best scientific theories.
- (P2) Mathematical objects play an indispensable explanatory role in science.
- (C) Hence, we ought rationally to believe in the existence of mathematical objects. [Baker 2009: 613]

Baker illustrates (2) with an elegant example. Biologists sought to explain why the time spent by the North American cicada underground in larval form before emerging as an adult is a prime number of years (13 or 17) rather than some other number, subject to the biological constraints of that species.²²

Baker reconstructs the explanation thus:

- (1) Having a life-cycle period which minimizes intersection with other (nearby / lower) periods is evolutionarily advantageous. [**biological 'law'**]²³
- (2) Prime periods minimize intersection (compared to non-prime periods). [**number theoretic theorem**]
- (3) Hence organisms with periodic life-cycles are likely to evolve periods that are prime. [**'mixed' biological / mathematical law**] [Baker 2005: 233]

²⁰See also Lyon and Colyvan [2008] where the authors argue that, even if phase-space theories can be nominalized, they would lack the explanatory power of mathematicized phase-space theories.

²¹In fact, something very much akin (but subtly different) to Baker's enhanced indispensability argument is anticipated by Field [1989] who suggests an argument that an acceptance of inference to the best explanation commits one to the existence of mathematical objects:

[S]uppose (a) that we have certain beliefs, beliefs about 'the phenomena', which we are unwilling to give up; (b) that this class of 'phenomena' that we believe in is large and complex; (c) that we have a pretty good explanation of these phenomena (in the sense of, a relatively simple non-ad hoc body of principles from which they follow); and (d) one of the assumptions that appears in this explanation is claim S, and we are pretty sure that no explanation of the phenomena that does without claim S is possible. The claim of 'inference to the best explanation' is that under these circumstances we have a strong reason to believe claim S. [Field 1989: 15].

²²Although Baker need not have looked so far afield as entomology to furnish us with such an example. Stewart Shapiro once offered the following:

[S]ometimes explanations of physical phenomena involve mathematical facts. An explanation of why a package of 191 tiles will not cover a rectangular area (unless it is one tile wide) might mention the fact that 191 is a prime number. [Shapiro 2000: 217]

²³A word of explanation is merited here. Minimizing intersection with other (nearby / lower) periods is evolutionarily advantageous because minimizing overlap with other periodic creatures is beneficial when they are predators (for obvious reasons), but also when they are a different subspecies of cicada (because mating with another subspecies of cicada would produce offspring that would have different life-cycle periods from other cicadas and, hence, would not be available for breeding at the same time as other cicadas).

A number of responses have appeared rejecting (P2) (Melia [2002],²⁴ Bangu [2008], Daly and Langford [2009] and Saatsi [2011]). While we will not discuss these responses in any detail here, there are quite general reasons to think that any rejection of (P2) will be problematic. Without a worked-out and agreed-upon theory of explanation there may be no principled way in which to settle the issue. Often the charge levelled by those who dispute (P2) is that the role of mathematics is *representational* rather than explanatory *per se*. The strategy is this: if it can be shown that the disputed applications of mathematics are representational then the view that they are explanatory will be undermined. However, it is not clear that representational and explanatory roles are mutually exclusive. Baker, for instance, adopts a ‘pragmatic account’ of explanation, on which an explanation of a phenomenon is constituted by an answer to a why-question, which demonstrates why the phenomenon under examination is more likely than its alternatives. If we endorse this pragmatic account it is clear that Baker’s example *does* constitute a mathematical explanation, regardless of whether the role of mathematics is in this instance representational. Another way nominalists may try to motivate a rejection of (P2) is to offer alternative nominalistic explanations to replace current mathematical explanations. Saatsi [2011], for instance, suggests a replacement nominalistic explanation for (1)-(3), involving a fact about time—There is a unique intersection minimizing period T_x for periods in the range $[T_1, \dots, T_2]$ years—replacing the fact (2) about prime numbers. Here, we are left with *two* explanations, one mathematical and one non-mathematical. Saatsi argues that the mathematical explanation is not more explanatory than its nominalistic counterpart and, as such, the nominalistic version should be favoured. Yet, while there is much to be gleaned from his discussion, it is again far from obvious that there is a principled way to determine whether mathematical explanations are or are not more fundamental than nominalistic ones. A debate that continued along this path is in danger of being dragged into that philosophical slough; the internecine trading of intuitions. Better then for the nominalist to sever the beast at the head, provided there are grounds for doing so. A more fundamental objection however is that whether or not these mathematical explanations are, in the end, indispensable, is an irrelevance. The considerations here parallel exactly those when we considered the role of indispensability itself in a defence of nominalism and found it to be strictly irrelevant. Let us imagine that for every mathematical explanation there is a replacement nominalistic one. Are the former any less explanations for our elision of them? Surely not! If there are mathematical explanations in science then we need to establish whether they warrant belief in the existence of mathematical objects. If they do then they will continue to do so whether or not we ‘dispense with them. If they do not, then their indispensability is irrelevant again. What is of interest to the ontologist is whether mathematics plays an explanatory role in science *simpliciter*, not whether that role turns out to be indispensable.²⁵

²⁴Strictly speaking of course, this is not a response, given that it precedes Baker’s paper.

²⁵It may be thought that nominalistic explanations, if available, would *supersede* mathematical ones, just as Darwinian explanations of organisms are thought to supersede Aristotelian ones. However, this kind of supersession requires an *incompatibility* between the competing explanations, and it isn’t the case that intrinsic nominalistic explanations are incompatible with extrinsic mathematical ones. As such, even if an intrinsic physical explanation is ‘better’ than a mathematical one, the mathematical explanation cannot be set aside or deemed incorrect for this reason.

Whereas attention has primarily been focussed on (P2), I will argue that (P1) should be rejected. This possibility has been largely ignored I think because making use of an explanation whilst denying the truth of the explicans can look like intellectual doublethink *in excelsis*. Baker for instance points out that while the following is explanatory:

- (17) The ‘missing’ mass-energy in the products of the electron decay is due to a neutrino having been emitted. [Baker 2009: 628]

if it were weakened into a subjunctive conditional, it would cease to be:

- (18) If there were neutrinos then the ‘missing’ mass-energy in the products of the electron decay would be due to a neutrino having been emitted. [ibid.]

This much is certainly true. To claim that ‘the ‘missing’ mass-energy in the products of the electron decay is due to a neutrino having been emitted, but neutrinos don’t exist’, is intellectual doublethink of the worst kind.²⁶ It is telling however that the example Baker uses is an instance of causal explanation. While it is obvious that to assert ‘*c* caused *e*’ whilst adding parenthetically that *c* does not exist, is incoherent, it is not obvious that denying the existence of something quantified over in an explanation, regardless of the *kind* of explanation, will be incoherent. In fact Baker’s case for the enhanced indispensability argument explicitly trades on a supposed parity between causal and mathematical explanations. He says of his example that it

parallels cases of explanation involving concrete theoretical posits, which are unproblematic common ground for both platonists and nominalists in the indispensability debate. Why is the light from certain distant galaxies getting bent? Because there is a black hole between us and the distant galaxies. Why do periodical cicadas have prime periods? Because prime numbers minimize their frequency of intersection with other period lengths. [Baker 2005: 235]

Contrary to the supposed parity, explanation is multifaceted and comes in many taxa: *Causal Explanation*; *Theoretical Reduction*; *Conceptual Explanation*; *Unification*; *Structural Explanation*; *Deductive-Nomological Explanation*; *Teleological Explanation*; *Statistical/Probabilistic Explanation*; no doubt there are other kinds of explanation besides these. My goal at this point is modest: it is simply to note that whilst there is surely something common to all these kinds of explanation—if only that they all play some kind of demystifying role—it most certainly isn’t clear that what applies to one form of explanation must apply to another.

So stand the contours of this debate; but the assumptions which lie just beneath its epidermic surface, and give it its peculiar shape, have not been made adequately explicit, if we are to be in a position to assess its outcome. Two fundamental questions raise themselves:

- (1) How, in the most general terms, is mathematics applied?

²⁶This I take it is why Melia’s [2000] strategy of ‘weaseling’, while touching upon *something* significant about the role of mathematics in science, has not commended itself to more people.

(2) Can applications of this sort be explanatory?

To answer (2), we must first answer a third subsidiary question:

(3) What constitutes an explanation?

Only then will we be in a position to answer, firstly, whether mathematics features in physical explanations, and, secondly, whether this would provide warrant that mathematical objects exist. We will turn firstly to the application of mathematics and its role in *modelling* the world. What I will say about this is not, I hope, particularly controversial. However, I think that the import of these uncontroversial truths has not been sufficiently appreciated with respect to the rather more contested issue of the relationship between mathematical explanation in the sciences and the existence of mathematical objects. Following this, I will suggest minimal conditions for explanatoriness. Whilst applications of mathematics can indeed be explanatory in this minimal sense, I will argue that applied mathematics is explanatory *in virtue* of its role in modelling the world and that *this* role is unproblematic for the nominalist. It is to these issues which we shall now turn.

5.3 Applying Mathematics

There have developed in the literature two strands of discussion on which we shall draw. One is that already discussed, issuing from debates and discussions directly pertaining to the philosophy of mathematics in which the application of mathematics and its various ontological implications (or lack thereof) is brought to the fore. The other strand pertains to the philosophy of science more generally, where a burgeoning literature on scientific representation and modelling has developed, often with an eye to the role of fictions in science. A representative, though not comprehensive, selection of theorists working in this field would include Gabriele Contessa [2007, 2010, 2011], Steffen Ducheyne [2006], Bas van Fraassen [2008] Steven French [2003, 2010, 2011], Roman Frigg [2006, 2010a, 2010b, 2010c] Ronald Giere [1999, 2004, forthcoming], Peter Godfrey-Smith [2009], Mauricio Suarez [1999a, 1999b, 2003, 2004, 2008, 2009a, 2009b, 2010a, 2010b] and Paul Teller [2001]. This second strand contains insights that the philosopher of mathematics should heed. Here I will try to offer something of a synthesis of these two lines of thought, that by their entwining we might unravel some of the issues about which there has been unclarity.

5.3.1 How Scientists Represent the World

The philosophical literature on scientific representation has prevalingly concerned itself with the topic of scientific *modelling*. This of course is a central part of the scientific enterprise, but it should be acknowledged that scientific representation is by no means limited to representation via models. This is a point that in recent years some authors have emphasized. Tarja Knuutila, for instance, has noted:

It is somewhat paradoxical that most philosophical discussion on scientific representation so far has been conducted in the context of modeling, whereas scientific endeavor employs manifold representations that are not readily called models. Such representations include visual and graphic displays on paper and on screen such as pictures, photographs, audio-graphic and three-dimensional (3D) images, as well as chart recordings, numerical representations, tables, textual accounts, and symbolic renderings of scientific entities such as chemical formulas. In fact, modeling presupposes these more “basic” forms of representation, which organize the data and present theoretical objects and fields of interest in some representational medium. [Knuutila 2007: 217]

In a similar vein Peter Godfrey-Smith [2006] criticizes the “semantic” view of theories. According to the view, scientific theories can be identified with the class of structures, in the set-theoretic sense of that term, which are its models. A structure $S = [D, R]$ consists of a non-empty set D of objects, and a non-empty set R of relations on D . Scientists construct models, which are satisfied, or fail to be satisfied, by target physical systems.²⁷ What Godfrey-Smith takes exception to is the claim that this characterizes *all* scientific theorizing; a claim that can be disabused by a survey of the scientific literature.²⁸ A similar point is also made by Steven French [2010] regarding the debate over whether philosophers of science should adopt the semantic view of theories, or the “syntactic” one, that takes theories to be linguistic entities. As French notes, ‘a variety of representational resources are available and even if there appear to be decisive reasons for preferring one such resource over others, such reasons might be viewed as representationally pragmatic, as far as the philosophy of science is concerned.’ [ibid.:235] There are myriad means by which scientists represent the world; there is no profit in attempting to artificially subsume all scientific representation under one method.

The notion of a model itself, and the role models play in scientific theorizing and representation, has been finessed in recent discussions. Not all models of course are mathematical; engineers may use a scale model of an aeroplane wing in a wind tunnel to represent a full-scale counterpart, or biologists idealized causal mechanisms to explain how traits may have evolved in an organism.²⁹ Scientific models are then quite diverse. Moreover, the models actually employed by scientists do not inertly reflect the physical world. Paul Teller [2001] urges philosophers to jettison the ‘perfect model model’. Models are not perfect mirror images of the systems they represent:

[T]he messiness of initial conditions, theoretical virtues such as simplicity, and other such constraints on theorizing are really matters which are relative to our intellectual capacities. [...] To have theories which we can actually apply in describing and understanding the

²⁷This is a view of scientific models that can be traced back to Suppes [1960, 1969].

²⁸Godfrey-Smith takes an example from evolutionary biology, Leo Buss’ *The Evolution of Individuality*, as a case in point, noting:

There are no formal mathematical methods in Buss’ book. And further, there are no overt models of any other kind. Buss’ entire argument is based on the causal roles and consequences of actual cellular machineries, actual environmental circumstances, and actual developmental sequences. [ibid.: 732]

²⁹As van Fraassen notes, “‘Model’ is a metaphor, whose base is the simply constructed table-top model. We use this metaphor when we talk of cosmological models, Hilbert space models and the like. We could have used the word “map”, and made maps the base of our metaphor equally well.’ [van Fraassen 2008: 83].

world we have no choice but to work with nature to do what it does not sufficiently do by itself: We must simplify further. [Teller 2001: 394]

This approximation, simplification, idealization and distortion is what motivates Nancy Cartwright [1983] to provocatively and notoriously call laws of physics ‘lies’. Yet to speak in such terms can be infelicitous, since scientific representation involves *deliberate* approximation, simplification idealization and distortion; it is not the goal of science to be a mirror of nature. This is a point that Catherine Elgin has stressed:

These would be embarrassing admissions if models were supposed to accurately reflect the facts. But they are not. Science is not, cannot be, and ought not be, a mirror of nature. Rather, science embodies an understanding of nature. Since understanding is not mirroring, failures of mirroring are not necessarily failures of understanding. [Elgin 2007: 77]

Scientific representations, Elgin contends, provide understanding not in virtue of mirroring the behaviour of concrete systems but by *exemplifying* salient features of them and by that means providing epistemic access to otherwise obscured properties of those systems. There is a difference between an object merely instantiating a property and exemplifying a property.³⁰ A Farrow and Ball colour card of ‘antique white’ instantiates many properties (a certain mass, shape etc.) but exemplifies the colour antique white. In this case the property exemplified by the colour card is not metaphysically privileged over its other properties; it exemplifies antique white because that is its *role* in our practices, it is *used* to display that property. Idealizations are indispensable in scientific representation because they exemplify properties that perfect mirrors of the target systems would not. It may be, to use Elgin’s own example, that every real swinging bob is subject to friction, yet by modelling idealized pendulums that are not, we can disentangle distinct properties of pendular systems, and by this means see how they function and why they behave as they do, in a way that would be impossible without idealization of some sort. In systems where disparate properties come bundled together in fact, idealizations are indispensable for a scientific understanding of their behaviour. The situation with regards to the construction of models is akin to that of the laboratory. In laboratories artificially controlled environments are constructed in order to isolate systems from external influences and interactions, in order that the systems under scrutiny can exemplify special causal relationships.³¹ We may never observe these causal relations (or the counterfactual relationship they necessitate) in nature because they are obscured by other factors. Yet, the relation is *instantiated* in nature even though it is not there *exemplified*. The role of models is analogous to that of laboratories; both actively employ a degree of artificiality in order to afford epistemic access to nature in a way that would be impossible through passive observation or mirroring.

³⁰In taking in terms of ‘instantiation’ I do not intend to invoke the existence of abstract properties. Nominalists do not believe in universals (although there is logical space to be a *mathematical* nominalist who does). The nominalist holds that properties do not exist, but she does not hold the crazy view that objects do not have properties, rather that there is no domain of free-floating properties that exist independent of any object that instantiates them. In claiming that ‘*o* instantiates *P*’, I mean only that ‘*Po*’.

³¹Cf. Rouse [1987, 2007] for an interesting discussion of this role of laboratories.

5.3.2 The Possibility of Applied Mathematics

Scientific representation then is multifarious. Modelling, as a subset of these practices, is itself a diverse affair and belies the view of scientific representation as a mirror of nature. Although our interest is in *mathematical* modelling, it will be instructive to bear in mind the lessons of scientific modelling and representation more generally. Before we turn to the main theories of applied mathematics in the literature however, it is worth addressing a more foundational question: how it is that mathematics *could* be applied in the first place. There is, as many have noted, an intimate connexion between mathematics and physics, yet the two are essentially distinct:

The *physical object* cannot be determined by axioms and definitions. It is a thing of the real world, not an object of the logical world of mathematics. Offhand it looks as if the method of representing physical events by mathematical equations is the same as that of mathematics. Physics has developed the method of defining one magnitude in terms of others by relating them to more and more general magnitudes and by ultimately arriving at “axioms”, that is, the fundamental equations of physics. Yet what is obtained in this fashion is just a system of mathematical relations. What is lacking in such a system is a statement regarding the significance of physics, the assertion that the system of equations is *true for reality*. [Reichenbach 1965: 36]

We must cash out what is involved in the ‘statement regarding the significance of physics, the assertion that the system of equations is *true for reality*’. Mathematics is applied to the concrete domain via theoretical hypotheses which, as Ronald Giere briskly puts it, ‘have the form: ‘Such and such real system is a system of the type defined by the theory’’ [Giere 1979: 70]. How is this act of coördination between the mathematical domain and the concrete possible, how can scientists come to represent the world mathematically? One is tempted here, as van Fraassen puts it, to ‘impose a parallel vocabulary and declare victory’ [van Fraassen 2008: 138], and make a claim of the following sort:

[A] point in Minkowski space corresponds to a real or physical spacetime point, which is a compact convex part of the WORLD with zero measure. [ibid.]

But, as van Fraassen points out, to describe the application of mathematics in this parlance is only meaningful if one *already* understands how the concrete world can be described mathematically. This description depicts the WORLD in mathematical terms, as a model in the set-theoretic sense, with a domain of objects and a set of relations on those objects. That is to say, it *presupposes* an ability to represent the concrete domain mathematically rather than making explicit how descriptions of this sort could take place. The heavy-duty platonist may be tempted to go a step further here and say that what this shows us is that the WORLD *is* a mathematical structure. If we are to understand the applicability of mathematics to the *concrete* world then we are charged with finding a more explanatory (and nominalistically acceptable) account of how this could take place. The route to understanding mathematical representation is through measurement.³² Central to van Fraassen’s account of measurement is

³²I follow van Fraassen on this point: ‘[A] theory would remain a piece of pure mathematics, and not an empirical theory at all, if its terms were not linked to measurement procedures.’ [van Fraassen 2008: 115]

the notion of a *logical space*: *The act of measurement is an act—performed in accordance with certain operational rules—of locating an item in a logical space.* [van Fraassen 2008: 165]³³

Agents represent a given logical space, for instance \mathbb{R} . The act of measuring an object o , is assigning o a position in this logical space, based on measurement outcomes. This involves physical interactions between agents and the world: measurement outcomes are concrete events or objects; end states of measurement apparatus, images on screens hooked up to fMRI scanners, photographs, charts, lists of numbers and the like. In assigning o a position in \mathbb{R} , the agents license *reasoning* about o in particular ways, specified by the mathematical component of the theory. As such, measurement is a kind of representation: objects are represented *as* thus or so. What is measured, and hence given a location in a logical space, is not defined independently of this act of location, but rather is defined by this act of location. This is a way of saying that measurement is *theory-laden*.

There is a degree of convention in which measurement scales one chooses to adopt. While it may be practically felicitous to, say, talk in terms of Fahrenheit rather than Celsius in certain social contexts, one scale is not more correct than the other. These two distinct scales provide two distinct means by which one can represent the facts about temperature, but in both cases the same facts are represented. Given that the choice of scale is arbitrary, the *significance* of these scales cannot reside in the particular numbers they assign to temperatures. What *is* significant about a measurement is that which remains invariant under transformations of the scale. In the case of temperature, as well as many other measurement scales, the more fundamental logical space in which an object or event is located then is a more abstract structure characterised by the set of numerical transformations under which a kind of measurement is invariant. One can distinguish different kinds of scales by the kinds of transformations under which they are invariant. By this method, five different kinds of scale are usually distinguished

1. **Absolute** (e.g. counting): Absolute scales are invariant under *identity transformations*. f is an identity transformation iff for every real number x in the domain of f , $f(x) = x$.
2. **Ratio** (e.g. mass, Kelvin temperature scale): Ratio scales are invariant under *similarity transformations*. f is a similarity transformation iff there exists a positive number α such that for every real number x in the domain of f , $f(x) = \alpha x$.
3. **Interval** (e.g. Celsius and Fahrenheit temperature scales): Interval scales are invariant under *linear transformations*. f is a linear transformation iff there is a positive number α and a number β such that for every number x in the domain of f , $f(x) = \alpha x + \beta$.
4. **Ordinal** (e.g. Beaufort wind scale, Mohs hardness scale): Ordinal scales are invariant under *monotone transformations*. f is a monotone-increasing transformation iff for every x, y in the domain of f , if $x < y$ then $f(x) < f(y)$. f is a monotone-decreasing transformation iff for every

³³See also Davis Baird who talks of a ‘field of possibilities’ in place of van Fraassen’s ‘logical space’. He states: ‘Measurement presupposes representation, for measuring something locates it in an ordered space of possible measurement outcomes.’ [Baird 2004: 12, quoted in van Fraassen 2008: 165]

x, y in the domain of f , if $x < y$ then $f(x) > f(y)$. f is a monotone transformation iff f is a monotone-increasing transformation or a monotone-decreasing transformation.

5. **Nominal** (e.g. players' numbers in a football team): Nominal scales are invariant under *one-to-one functions*. f is a one-to-one function iff for every x, y in the domain of f , if $f(x) = f(y)$ then $x = y$.

Celsius and Fahrenheit scales are two coordinate systems that both pick out the same *logical space*; a space which is determined by its more abstract relational structure. Determining what transformations are permissible—*viz.* what features of the measured system are to be captured in the measurement practice—is a practical matter involving the purposes of those measuring the system and what relationships it is practically possible to access through measuring apparatus. As van Fraassen notes, these are factors central to an account of measurement that lie outwith the province of a purely mathematical theory of measurement:

[W]hat decides what is and is not important, what are precisely the relevant features to be preserved in all and only the admissible transformations of the scale, is not decided by mathematical characteristics. It is decided by two factors: the measurer's purpose and the relationships that can be assessed instrumentally or practically on which to base the assignment. These factors are outside the domain of the mathematical theory of measurement. [van Fraassen 2008: 161]

Adopting van Fraassen's example, a simple theory of gases which describes their properties along three parameters, P(ressure), V(olume) and T(emperature), represents a three-dimensional mathematical space. The goal of measurement procedures is to locate measured bodies of gas in this logical space, \mathbb{R}^3 . While the reals or ordered n-tuples of the reals are but one, albeit paradigmatic, logical space that agents represent in order to carry our measurements; the broader notion of a logical space is ubiquitous in scientific measurement: 'The HSB color space, with dimensions hue, brightness, and saturation is a good example of a logical space, but so is the PVT space in elementary gas theory, phase space in classical mechanics, Hilbert space in quantum mechanics; space and time themselves also serve as examples.' [van Fraassen 2008: 164]³⁴

What this draws to the fore is the inadequacy of a methodologically platonistic solution to the problem of coordination. We cannot understand the distinction, and the relationship, between pure mathematics and applied mathematics by positing more metaphysical relations between mathematical and concrete objects. What is called for is an explanation of a methodologically pragmatist sort. This involves not only physical interactions between agents and measurement outcomes, but also cannot be understood in the absence of an account of the how these worldly interactions are embedded in a particular constellation of practices. These are the factors that lie outwith a purely mathematical account of applied mathematics.

³⁴Even this picture is something of an idealization. Models are not subject to *direct* empirical testing, but through the intermediary of *data models*.

5.3.3 Theories of Applied Mathematics in the Literature

Our discussion of applied mathematics is motivated in part by a need to get clear on the hazy issue of whether mathematics can be involved in explanations of physical phenomenon. This discussion however can serve a dual purpose, for not only does the application of mathematics underly the sort of mathematical explanations the Enhanced Indispensability Argument seeks to capitalize on, but the very applicability of mathematics *itself* is often thought to be somehow inexplicable or problematic given the supposition of nominalism. Field, for instance, states:

It has been widely held that the most serious difficulty facing any non-realist philosophy of mathematics is *the problem of application*: how can one account for the utility that reasoning about mathematical entities has for disciplines other than mathematics, if mathematics isn't construed in a realist fashion? Applied to deflationism, the problem is: how can one explain the applicability of mathematics to disciplines other than mathematics, without assuming that ordinary mathematical claims (including those claims that assert the existence of mathematical entities) are true? [Field 1989: 94]

It has not been so widely considered that explaining the applicability of mathematics may be a problem *per se*, as apposed to merely a problem for nominalists of various stripes. Shapiro [1983 and 1997] is a notable exception to this trend, as is Balaguer [1996 and 1998], and quite rightly so; for the mere truth of mathematics would certainly not be sufficient in itself to explain its successful application. Astrologers may make true claims about the perambulations of bodies in our solar system, but this truth would not in itself be sufficient to explain why such facts are relevant to our love lives. So another problem becomes apparent: how can one explain the applicability of mathematics to empirical science *given* the existence of mathematical objects? And with this in mind we begin to wonder whether explaining the applicability of mathematics is more difficult for the nominalist than it is for the realist, or indeed whether they may both be able to explain this applicability in the same way.

The Mapping Account

The primary modes of mathematical application can be divided into four main categories³⁵:

1. **Arithmetic**

Counting (assigning numbers to concepts of classes of concreta);

2. **Analysis**

(a) *Measurement* (introducing certain real-valued functions from concreta to real numbers: e.g. measurement scales, real-valued classical fields on space-time);

(b) *Geometry* (introducing co-ordinate systems on space time);

3. **Set Theory**

³⁵The taxonomy is based on that of Ketland [1998: 44].

Forming collections of concreta (sets and relations-in-extension);

4. Exemplification

Assigning structures as “representations” of concrete systems.

In recent years attention has been given to describing a more general account which encompasses these four modes of application. Mathematics is a wellspring of structures, and this fact has been thought to account for its uncanny utility in empirical science. Apropos of this, Christopher Pincock [2004 and 2007] has proffered a structural account of the application of mathematics which he refers to as the ‘mapping account’, and other theorists (Balaguer 1998: 144, Baker 2003, Leng 2002 and Azzouni 2004) have alluded to the application of mathematics in a way that suggests they may also have some kind of structural account in mind³⁶. Here I will outline the first the structural account and then the ‘inferential conception of applied mathematics’ put forward by (the nominalist) Otávio Bueno and (the platonist) Mark Colyvan, which is essentially an extension of Pincock’s structural account.

The fundamental and simple notion underlying the structural account of the applications of mathematics is that we can produce mappings between those parts of the physical world which we wish to describe and various mathematical entities which preserve structure between that part of the physical world and the mathematical entity used to describe it. It is often not the case that there is a straightforward isomorphism between the relevant physical and mathematical entities; there can be more structure in the world than in the mathematics used to represent it, and, contrariwise, there can be more structure in the mathematics than in the physical systems it is used to represent. With this in mind, the mappings used are usually homomorphisms, epimorphisms and monomorphisms. For some language \mathcal{L} , an \mathcal{L} -structure consists of a domain D , with a function assigning to each constant symbol c in \mathcal{L} an element of D , to each n -ary function symbol F in \mathcal{L} an n -ary function on D , and for each n -ary relation symbol R in \mathcal{L} an n -ary relation on D . Let N and M be \mathcal{L} -structures. If N is homomorphic to M then there is a mapping $f: N \mapsto M$ such that for all c , $f(c_N) = c_M$, for all F , $f(F_N(a_1, \dots, a_n)) = F_M(f(a_1), \dots, f(a_n))$, and for all R , $(R_N(a_1, \dots, a_n)) = R_M(f(a_1), \dots, f(a_n))$. An epimorphism is a surjective homomorphism and a monomorphism is an injective homomorphism. Which mapping is used will depend on the relative structural complexity of N and M . If, for example, N has more structure than M , then the mapping could be either a homomorphism or an epimorphism. If M has more structure then the mapping could also be a monomorphism.

Although it tells us something useful about what must take place for mathematics to be applied, the existence of a mapping from a mathematical structure to a physical structure is, on its own, insufficient to make sense of the applicability of mathematics.³⁷ Bueno and Colyvan identify three shortcomings of the mapping account. Firstly it does not tell us what parts of the mathematical structure represent

³⁶In fact all of this theorizing is reminiscent of earlier technical work by Krantz, David H., R. Duncan Luce, Patrick Suppes, and Amos Tversky [1971].

³⁷Bueno and Colyvan construe Pincock as holding that ‘the existence of an appropriate mapping from a mathematical structure to a physical structure is sufficient to fully explain the particular application of the mathematical structure in question’ [ibid., p.346], although it is not clear that Pincock actually makes this strong claim.

features of the concrete system it is used to describe. Some interesting philosophical questions arise when M has more structure than the empirical domain, leaving some mathematical structure dangling free of obvious interpretation. Bueno and Colyvan give the example of calculating the landing position of a projectile whose initial velocity and position are known and whose only acceleration is due to gravity:

the displacement function for such a projectile is a quadratic with two real solutions—only one of which is physically significant. The problem is how do we know which parts of the mathematical structure represent and which parts do not? The mapping account is silent on this issue. [Bueno and Colyvan 2011: 349]

This is not to say that the mapping account is not viable, only that it is not a comprehensive account of the way mathematics is applied insofar as it does nothing to explicate the significance of cases such as these; something which a comprehensive account ought to. It is interesting to note that sometimes the dangling mathematical structure can be thought to have no physical interpretation when in fact it does. Lorentz took time dilation in his theory of length contraction to be a merely mathematical anomaly, but, post-relativity, scientists believe that this mathematical structure corresponds to real facts about time. The second problem Bueno and Colyvan see in the mapping account is an inability to explain the applicability of idealized mathematical models—for instance where fluids are represented as continuous, or springs as simple harmonic oscillators—where it is known that the model does not accurately represent the world. In these cases there is no mapping from the concrete system to the mathematical structure. Thirdly and finally, the mapping account is thought to be unable to explain the possibility of mathematical explanations of physical phenomena. According to Bueno and Colyvan: ‘The problem is simply that it is hard to see how a mere representational system can provide explanations and yet that is the only role mathematics is allowed to play in the mapping account’. [ibid.: 351]

5.3.4 The Inferential Account

The inferential account builds on the minimal analysis of the mapping account and focusses, as the name may suggest, on the fact that such mappings enable inferences to be made in the mathematical codomain whose results can then be interpreted as physical phenomena. Bueno and Colyvan usefully summarise the inferential account thus:

The crucial feature of the proposal is that the fundamental role of applied mathematics is inferential: by embedding certain features of the empirical world into a mathematical structure, it is possible to obtain inferences that would otherwise be extraordinarily hard (if not impossible) to obtain. [Bueno and Colvan 2011: 352]³⁸

³⁸This bears a striking resemblance to Suárez’s [2004] *Inferential Conception of Scientific Representation*, although transposed into a mathematical key, but we shall see that there are some important differences.

Also of import is the introduction of pragmatic considerations; the inferential account ‘is not purely structural, since it makes room for additional pragmatic and context-dependent features in the process of applying mathematics’ [ibid.] Of course there are roles for applied mathematics that would not immediately be described as ‘inferential’: unifying separate scientific theories³⁹; making novel predictions; and, more controversially, producing explanations of empirical phenomena (such as the previously discussed example of periodic cicadas). However, Bueno and Colyvan go on to say that: ‘All of these roles ... are ultimately tied to the ability to establish *inferential relations* between empirical phenomena and mathematical structures or among mathematical structures themselves.’ [ibid.]

How then are these inferences made? The inferential process involves three stages: *immersion*, *derivation* and *interpretation*:

Immersion: A mapping is made from the empirical set up to a mathematical structure. Features of the physical world are related to their counterparts in the mathematical structure.⁴⁰ The kind of mapping used (isomorphism, homomorphism, epimorphism or monomorphism) is dependent on the particular application in hand.

Derivation: Consequences are drawn in the mathematical structure being used.

Interpretation: The consequences drawn at the derivation stage are interpreted empirically. To do this, a mapping from the mathematics back to the empirical events or entities must be established. The kind of mapping used need not necessarily be of the same kind as in the immersion stage.

This conception of the application of mathematics bears a strong resemblance to the theory of scientific representation propounded by R.I.G. Hughes [1997, 2010], in which representation in the sciences involves denotation, demonstration and interpretation (DDI). Bueno and Colyvan see theirs as an extension of Hughes’ account, tailored to the more specific task of explaining *mathematical* representation in the sciences. To this end, they have replaced Hughes’ more generically referential notion of denotation with that of *interpretation*.⁴¹

Bueno and Colyvan distinguish between the empirical set up and the description of the empirical set up. In particular:

³⁹Colyvan [2002] gives an example of complex analysis being used to unify the mathematical theory of differential equations and thus to unify physical theories which employ differential equations.

⁴⁰The immersion step can apply to the application of mathematics to other parts of mathematics as well as to the application of mathematics to the physical world. Arithmetic can be ‘immersed’ in set theory, for example, in order to make inferences about the former.

⁴¹In fact, this understanding of the application of mathematics is not a new one. For instance, in Lin and Segel’s textbook *Mathematics Applied to Deterministic Problems in the Natural Sciences* we find the following description:

The process of using mathematics for increasing scientific understanding can be conveniently divided into the following three steps:

- (i) The *formulation* of the scientific problem in mathematical terms.
- (ii) The *solution* of the mathematical problems thus created.
- (iii) The *interpretation* of the solution and its empirical verification in scientific terms. [Lin and Segel 1988: 5, quoted in Pincock 2009: 176]

[T]here need [not] be a mathematics-free description of the empirical set up. Very often the only description of the set up available will invoke a great deal of mathematics. Thus, it will be hard even to talk about the empirical set up in question without leaning heavily on the mathematical structure, prior to the immersion step. The empirical set up is the relevant bits of the empirical world, not a mathematics-free description of it.' [ibid.: 354]⁴²

Mathematical posits often proxy for empirical objects that we could not describe without speaking in mathematical parlance.

In order to accommodate the fact that there are not *complete* mappings between the concrete world and mathematical structures, Bueno and Colyvan appeal to a formal framework that deals in partial structures. A partial structure is a framework $\langle D, R_i \rangle_{i \in I}$ where D is a non-empty set and each R_i is a partial relation; i.e. the partial structure defines for each R_i a set of n -tuples R_1 that satisfy R_i , a set of n -tuples R_2 that do not satisfy R_i and a set of n -tuples R_3 such that it is undefined whether or not they satisfy R_i , and where $R_1 \cap R_2 \cap R_3 = \emptyset$ and where $R_1 \cup R_2 \cup R_3 = D^n$. If we take two partial structures $A = \langle D, R_i \rangle_{i \in I}$ and $B = \langle D', R'_i \rangle_{i \in I}$ then a partial function $f : D \mapsto D'$ is a partial isomorphism between A and B if f preserves all the structure between A and B , i.e. if f is bijective and for all x and y in D , $R_{i1}(x, y) \leftrightarrow R'_{i1}(f(x), f(y))$ and $R_{i2}(x, y) \leftrightarrow R'_{i2}(f(x), f(y))$. Partial homomorphisms can be similarly defined: a partial function $f : D \mapsto D'$ is a partial homomorphism from A to B if for all x and y in D $R_{i1}(x, y) \rightarrow R'_{i1}(f(x), f(y))$ and $R_{i2}(x, y) \rightarrow R'_{i2}(f(x), f(y))$. In cases where R_3 and R'_3 are empty then the partial isomorphism or homomorphism is also a standard isomorphism or homomorphism.

The partial structures approach allows Bueno and Colyvan to accommodate certain kinds of idealizations that are common when mathematical structures are used to represent concrete systems. For instance, in their example of neoclassical economics, economic agents are represented as being ideally rational, in the economist's sense, such that they always maximize their profits. Since real economic agents often lack the information, the wherewithal or the will to maximize their profits, there cannot be *full* mappings between the concrete system of real people being modelled and the mathematical model; yet a partial mapping may obtain. Here the partial function is from the empirical set up to the mathematical model, but partial mappings can be made use of in the *interpretation* as well as the *immersion* stage. Here we find one point where the literature on the application of mathematics and the literature on scientific modelling connect: Bueno and French [2011] use the partial structures approach (developed previously by French and Ladyman [1999], Bueno, French and Ladyman [2002], da Costa and French [2003]) to ground a formal account of how scientific theories represent.

⁴²Compare Azzouni:

Quantification over mathematical entities is indispensably assertorically utilized to represent empirical phenomena. Given that geometrical entities are mathematical entities, [this holds] even for ordinary macro-objects. We represent the movements of such objects by imposing a space-time matrix on our pre-scientific macro-object descriptions. This enables the representation of shapes, distances, velocities, mass distribution, etc., of macro-objects to any realizable accuracy. [Azzouni 2009: 63]

5.3.5 Issues

The inferential account is an advance on the mapping account, not only because of the work done by the formal framework of partial structures but also because it ‘makes room for additional pragmatic and context-dependent features in the process of applying mathematics’ [Bueno and Colyvan 2011: 352], the import of which we will soon discuss. Like the bare mapping account, it surely contains important insights as to how mathematics is applied, but, also like the bare mapping account, it has limitations.

Batterman’s Critique

The inferential conception is itself a mapping account of mathematical explanations; it is not a rejection of the mapping account, but an embellishment of it. At the heart of mapping accounts is the claim that mathematics is applicable in virtue of there *existing* mappings from concrete entities to mathematical entities; it is possible to use mathematical language to describe or explain the behaviour of concrete systems because of the existence of these mappings. Robert Batterman has noted that not all applications of mathematics can be characterized in terms of the existence of such mappings:

In contrast to explanations that appeal to mathematical entities (or properties of such entities) to explain physical phenomena, there are explanations⁴³ which, while mathematical, do not make reference to such objects. Rather, these explanations appeal to (or better ‘involve’) mathematical *operations*. [Batterman 2010: 4]

In particular, Batterman has in mind operations that involve singular limits, i.e. those in which there is a qualitative *discontinuity* in the mathematics on either side of the limit point; where ‘the behavior as one approaches the limit is qualitatively different from the behavior one would have at the limit’ [ibid.: 17]. Consider the two quadratic equations

$$(1) \quad x^2 - \epsilon 2x - 2 = 0$$

$$(2) \quad \epsilon x^2 - 2x - 2 = 0$$

where ϵ is a parameter.⁴⁴ With (1) for any nonzero value of ϵ as it approaches the limit $\epsilon \rightarrow 0$ the equation has two solutions. Similarly, when $\epsilon = 0$ there remain two solutions, $+\sqrt{2}$ and $-\sqrt{2}$. (2) on the other hand behaves differently. As ϵ approaches 0 there are two solutions for the equation, yet when $\epsilon = 0$ there is one: $x = -1$. (2) contains a singularity; there is no sense in which the solutions as ϵ tends to zero can be measured as *close* to the solution when $\epsilon = 0$. As Batterman has emphasized in a number of places [2002, 2005, 2010] operations with this character play an important role in much scientific theorizing. He gives the example from condensed matter physics of the way in which different

⁴³Batterman is concerned to establish instances of mathematical *explanations* involving operations, but we need not (yet) enter into the debate over whether these are authentically explanatory. All that matters for our current purposes is whether there are successful *applications* of mathematics that cannot be characterized in terms of mappings.

⁴⁴The example is Batterman’s.

‘fluids’ all behave in the same way with respect to certain critical points. Fluids can be in solid, liquid and vapour states, and which state they are in is a function of both the temperature and the pressure of the fluid. In a system consisting of a fluid in a container, in which transitions take place from one state to another, such as when water is boiled in a kettle, liquid and vapour states of the liquid coexist. But there is a critical point C , with a critical temperature T_c and a critical pressure P_c which marks a qualitative change in the way the system behaves. Below T_c and P_c the fluid exists as a vapour, and in a system below T_c the fluid cannot make the transition from a vapour to a liquid without there being a period in which both vapour and liquid coexist in the container. In a system where the temperature is greater than T_c and the pressure greater than P_c on the other hand, it is possible (indeed actually the case), as the temperature decreases, for the fluid to change from its vapour state to a liquid state without there being any point at which both vapour and liquid are present in the container. This can be characterized by the ‘order parameter’ Ψ , whose value is given by the difference between the liquid density ρ_{liq} and the vapour density ρ_{vap} in the container:

$$\Psi = |\rho_{liq} - \rho_{vap}|$$

At the critical temperature T_c , ρ_{liq} and ρ_{vap} are such that Ψ is zero, whilst below T_c Ψ is nonzero. This delimits the contours of the ‘coexistence curve’, which marks where the system behaves qualitatively differently: when the temperature of the system moves above T_c , the order parameter Ψ is zero. It is a curious fact that, close to their critical temperature T_c , different fluids all have coexistence curves of the same shape. The ‘reduced temperature’ t is a measure of how far the system is from a critical point and is given by

$$t = \left| \frac{T - T_c}{T_c} \right|$$

t allows for the comparison of different fluids with different critical temperatures. The fact that all fluids share this feature means that their behaviour can be universally characterized by the equation:

$$\Psi = |\rho_{liq} - \rho_{vap}| \propto t^\beta$$

Here the value of β determines what shape the order parameter Ψ takes. One issue that requires explanation is why the shape of the order parameters Ψ for distinct fluids vary proportionally to a power law t^β . In providing an explanation, limits come into the picture. Batterman explains:

[O]ne essential feature of the explanation provided is the invocation of the so-called thermodynamic limit. This is the limit in which (roughly speaking) the number of particles of the system, e.g., the number of H₂O molecules in the tea kettle, approaches infinity. And, of course, this is an idealization: water in real tea kettles consists of a finite number of molecules. This limiting idealization is essential for the explanation because for a finite number of particles the statistical mechanical analogs of the thermodynamic functions cannot exhibit the nonanalytic behavior necessary to represent the qualitatively distinct behaviors we observe. [ibid., p7]

The model requires an infinite system, since without such no phase transitions can take place. Because of this, there is no mapping from the entities in the mathematical model to counterparts in the physical system being modelled. As such, the applicability of these sorts of models cannot be subsumed under a mapping account. Something else is going on; there must be some other feature of these models, or (as I will suggest) some feature of the way in which we use them, that accounts for their applicability. Nor can the idealizations inherent in these models be done away with. One can distinguish between ‘Galilean’ and ‘non-Galilean’ idealizations. Galilean idealizations are those idealizations about which some de-idealizing account can be given. Models may treat extended concrete objects as unextended point masses or ignore other relevant features such as air resistance. In these cases a story can be told about how the model may be made more representative by factoring in those things that the model ignores for the sake of computational tractability. Non-Galilean idealizations admit no such de-idealizing procedure. Batterman’s case from condensed matter physics is one such example: were the limiting idealization done away with the model would not become more accurate (but less computationally tractable), instead it would cease to capture the nonanalytic behaviour of the system it represents.

In essence the problem is that whilst the partial structures approach copes well with partial mismatches between the model and the system being modelled, such as where detail is lost in the process, it sits uneasily with cases in which useful parts of the mathematical apparatus find no physical correlates whatsoever. In these cases the mathematical model is not an abstract mirror-image of the physical system it is representing, instead it is what one can *do* with the model that matters.

Pragmatic Objections

Perhaps a formal account of representation can be further ameliorated to cope with models which contain limiting idealizations; but there are objections of a different kind that apply to *any* formal account of scientific representation. The inferential account described the application of mathematics in terms of a structure-preserving mapping from one domain N onto a second domain M . As I have stated, this captures something about the applicability of mathematics, but it is a partial picture of what the application of mathematics involves. Applying mathematics is an *activity* and must be understood in those terms. Bueno and Colyvan’s inferential account of mathematical applications *is* an advance over the purely semantic account precisely because it takes up pragmatic issues, left out by mapping accounts, which are essential to understanding the application of mathematics. Not only does the inferential account leave out pragmatic ingredients, in the absence of which our account will be incomplete, it also is mistaken in requiring the *existence* of a mapping from a mathematical structure to a concrete one. The two points are related.

Until this point we have followed Bueno, Colyvan, French and Batterman by conducting our discussion in terms of how it is that mathematical models can represent physical systems. To do so is to speak as if the object of our investigation is a *dyadic relationship* between mathematical structures and physical

systems, and indeed this is how Bueno and French explicitly characterise mathematical representation. In attempting to account for mathematical representation in terms of the existence of a partial homomorphism between a mathematical structure and a physical system, Bueno and French's account implies that mathematical structures *themselves* represent in virtue of the existence of this structural relationship. Whilst these authors do not ignore pragmatic issues, they do not accord them a central place in understanding the nature of representation. This I think is a mistake. Earlier in the chapter we noted that two strands had developed in the literature on scientific representation. The second strand has, in recent years, sought an understanding of representation that places *use* and other *pragmatic* notions at the center of theorizing about representation. Representation takes place, according to these accounts, not because some privileged dyadic relation obtains between that which is represented and that which represents it, but because agents use objects in such a way that their *role* is that of a representation. Representation is to be understood in terms of representational *practices*.

Why are inherently metaphysical accounts of representation, placing at their centre a metaphysical relationship between the thing being represented and thing representing it, inadequate? Mauricio Suárez, in an insightful paper, levels some powerful objections against metaphysical theories of representation. His targets are similarity and isomorphism theories :

The *similarity conception of representation* [sim]: X represents Y if and only if X is similar to Y.

The *isomorphism conception of representation* [iso]: X represents Y is and only if the structure exemplified by X is isomorphic to the structure exemplified by Y. [Suárez 2003: 227]⁴⁵

French and Bueno's partial homomorphism theory can be thought of as a more sophisticated descendent of [iso], and whilst not all of Suárez's criticisms apply straightforwardly to it, they highlight insuperable problems for metaphysical theories of representation of any sort.

Suárez presents five arguments against [sim] and [iso]. Firstly, not all representational devices work by means of [sim] or [iso]. Inscriptions of mathematical equations are used to represent concrete systems, but are not similar to them in any way. Nor are they isomorphic. Whilst the abstract structure determined by the equation may be isomorphic to the concrete system, it is not this structure that is used to represent the concrete system. Scientists do not generally investigate the formal properties of the structures determined by the differential equations they employ. Rather, they find solutions to these equations and test them against measurable features of the phenomena they are investigating. Inscriptions of mathematical equations *themselves* can be used to represent concrete systems without a similarity or isomorphism obtaining. Bueno and French [2011] claim 'It cannot be the case that similarity in general is not necessary for scientific representation, since in the absence of such, the successful facilitation of the usual practices of interpretation and inference ... would be nothing short of a miracle.' [Bueno and French

⁴⁵Suárez uses 'A' and 'B' to denote, respectively, the source and target of representation. For the sake of continuity I have replaced these with 'X' and 'Y' in this and all other citations.

2011: 29]. But representing concrete systems with differential equations is just such a “miracle”; so too are cases, ubiquitous in scientific representation and representation more generally, in which natural language *sentences* represent parts of the world. To explain the movement of neurotransmitters by saying that ‘Calcium ions enter and bind with the protein embedded in the membrane of synaptic vesicles docked at the release zone [of the presynaptic membrane]’ [Carlson 1998, p46] is a clear case of scientific representation. This objection then stands, and carries over to Colyvan, Bueno and French’s inferential and partial morphism accounts.

Secondly, [sim] and [iso] do not have the logical properties of representation. Representation is not symmetric; the a source X does not represent its target Y in virtue of the fact that Y represents X . A doodle of Churchill represents the man, but the man does not thereby represent the doodle. Representation is also non-transitive. The guide to Edinburgh’s National Gallery represents Titian’s *Diana and Actaeon* by describing it; the painting itself represents a Roman mythological goddess, but the guide does not represent the goddess. In addition, representation is non-reflexive; in most circumstances things do not represent themselves. [iso] falls foul of all of these these logical properties. Whilst similarity is not transitive, it is symmetric and reflexive. Bueno and French note that homomorphisms are not symmetric and conclude that their approach ‘does not fall prey to the criticisms levelled at ‘similarity’ [and isomorphism] theories of representation’ [Bueno and French 2011: 5]. This conclusion ignores the fact that homomorphisms are both reflexive and transitive.

Thirdly, [sim] and [iso] cannot account for *misrepresentation*. Many representations are sufficiently inaccurate that conditions such as [sim] and [iso] could not hold for them. Newtonian mechanics misrepresents the solar system insofar as it is inaccurate with respect to the actual motion of the planets. Yet, Newtonian mechanics can be used to represent the solar system: ‘*Misrepresentation is a species of representation*’ [van Fraassen 2008: 15]. Bueno and French respond to this objection by claiming that when scientific representations are discovered to be inaccurate they are replaced by models of greater accuracy. The relationship between the new model and its discarded predecessor can itself then be characterised in terms of partial morphisms from the new model to the old. Three points are in order here. In the first place, this is not an accurate representation of all scientific practice. Many scientists choose to work with models that they *know* misrepresent their targets in some ways, but do so because these inaccurate models remain useful predictive or even explanatory tools. In the second place, it is not clear that any comparable procedure would be available in the case of non-Galilean idealizations such as those described by Batterman. Finally, the response is wide of the mark: even if partial morphisms can be used to articulate the relationships between more and less accurate representations, it remains the case that partial morphisms are neither necessary nor sufficient for scientific representation, on the grounds that scientific representations can be sufficiently inaccurate that no partial morphism obtains between the model and its target.

Fourthly, [sim] and [iso] are not necessary for representation. This is a point that has already been

addressed above in more specific forms, although more can be said, particularly in the case of artistic representation. David Hockney's *Swiss Landscape* depicts the Alps, but in such a rudimentary and distorted way that it is not similar or isomorphic to them.

Fifthly, [sim] and [iso] are not sufficient for representation. A snowflake falls in the Andes, and, by chance, an identical snowflake is simultaneously dropping over Aberdeen. The Aberdonian snowflake exemplifies a structure perfectly isomorphic to its Andean counterpart, yet is not a representation of it. Here too, the point carries over to the partial morphism account, although Bueno and French acknowledge that partial morphisms are not sufficient for representation.

What is interesting about these objections is not primarily that they present problems for a particular formal theory of representation—even a sophisticated theory such as Bueno and French's—but rather that they point towards a different way of understanding representation. It is not merely that, as is generally recognized, the formal theories of representation proposed thus far do not provide sufficient conditions for representation. What is more pressing is that it is difficult to see how they *could*, for any structure-preserving mapping will be reflexive, but it is a datum that not everything that represents, represents itself. In the same way, it is difficult to see how any formal conditions could be necessary for representation, because of the possibility of misrepresentation. The proponent of a formal theory of representation faces a dilemma. On the one hand, if the formal conditions are strong enough to be interesting they will not account for cases of misrepresentation. On the other hand, if the formal conditions are sufficiently weakened to take these cases on board (which is to say, weakened to the point of triviality), they will no longer be interesting, since they will obtain between the source and any number of targets which are not in fact represented by it. This suggests that pursuing theories of representation of this sort is somehow misguided, but how so? Because metaphysical theories of representation take what is fundamental to representation to be a dyadic relation between the source X and target Y of representation, they entail that knowledge of the intrinsic properties of X , Y and their relation is sufficient to determine whether X represents Y . This however is problematic. Bas van Fraassen suggests the following thought experiment:

Imagine: I have acquired a famous photograph of the Eiffel Tower, *Au Pont de l'Alma* by Doisneau. It hangs on my wall, but I scan it and print the scanned image. The print is an image too—*what does it represent?* The Eiffel Tower seen from the Pont de l'Alma, or the famous photograph?

There is no single immediately and obviously right answer; there couldn't be. It depends on what I do with the thing. [van Fraassen 2008: 21]

By insisting that some special kind of relation of representation must obtain between the source and the target of representation, metaphysical theories of representation flout the fact that representation can be, and frequently is, a matter of stipulation. I raise my dexter hand and say "This is Churchill", now my sinister hand, "This is Hitler", and a representational practice is already underway. (Almost?) anything can represent anything else, if one so chooses. These representations may or may not be felicitous—in

the case at hand, this will depend on what I choose to do with my hands; enact a scene perhaps—but they are representations nonetheless. Suárez makes a useful distinction between the *means* of representation and the *constituents* of representation.

The fact that we use a particular relation (say, similarity) between X and Y to, say, infer Y's properties by reasoning about X's properties, should not be taken to mean that this relation is what *constitutes* the representation by X of Y. [Suárez 2003: 230]

Fundamentally, metaphysical accounts of representation are mistaken because they misidentify the constituents of a representation for its means. The *raison d'être* of representation is surrogative reasoning: if X represents Y then one can infer features of Y by reasoning about X. In cases where X is similar or isomorphic to Y, this similarity or isomorphism can be made use of in order to carry out such surrogative reasoning; it is the *means* by which X can be used to (successfully) represent Y, but it is not *constitutive* of the representation; it is not the fact that makes X a representation of Y in the first place.

The representational 'force' of a given source—which target it represents—is fixed then by the intentional attitudes of those who employ it. So, in addition to avoiding the five objections Suárez levels at [sim] and [iso], a sound theory of representation must then account for the stipulative nature of representation. While the representational force and hence denotation of a source may be a matter of the intentional attitudes of those who employ it, denotation is not sufficient for scientific representation. Scientific representations are not merely representations; their purpose is to provide information regarding their targets. Suárez's own deflationary theory of representation [Suárez 2004] takes this role of scientific representations to be captured by their ability to facilitate the making of inferences by agents who make use of them. Suárez rejects the aim of finding a 'substantive' theory of representation, furnishing universal necessary and sufficient conditions for representation, but takes it that one can describe its most general features in an illuminating way. He suggests the following conditions:

[Inf] X represents Y only if (i) the representational force of X points towards Y, and (ii) X allows competent and informed agents to draw specific inferences regarding Y. [Suárez 2004: 773]

[Inf] avoids the five criticisms of [sim] and [iso]. Firstly, it allows that many different mediums can be used in scientific representation, and that these can facilitate inference-making without being similar to or bearing particular *structural* relationships to their targets. Straightforward similarity *could* of course be the vehicle by which inferences are made possible, but it is not supposed in advance that these are the *only* means by which this can take place; differential equations and natural language sentences are not wrongly excluded from scientific representations on this account. Secondly, [Inf] captures the logical properties of representation. It is nonreflexive, nonsymmetric and nontransitive, as a result of accommodating the stipulative, intentional nature of the direction of representation. Thirdly, misrepresentation is unproblematic on this account. Condition (ii) demands that the source allows competent and informed agents to draw conclusions about their target, but not that these conclusions are accurate. Fourthly, and relatedly, [Inf] provides plausibly necessary conditions for representation. Metaphysical

accounts of representation are too strong because if sources are required to, in some sense, mirror their targets, then cases of inaccurate representation ruled out. Finally, [Inf] is proffered as a set of necessary conditions for representation, so the objection that they are not sufficient is trivially ruled out. This said, the kind of examples that show that metaphysical accounts of representation fail to provide sufficient conditions—such as cases of isomorphic objects that do not represent each other—do not apply to [Inf], because condition (1) requires that a source can only represent its target if this representational force is so intended. [Inf] also explains the allure of the picture of science as a mirror of nature. A source that really was a mirror of its target would—setting aside issues of computational intractability and the like—if used appropriately, provide those who make use of it an inferential grasp of the target. The dual mistakes inherent in the picture of science as a mirror of nature are, on the one hand, the move from noting the sufficiency of a mirror (along with appropriate pragmatic conditions) to represent its target, to taking this to be a *necessary* condition for representation, and on the other hand to fixate on the source and target themselves, suppressing the pragmatic elements of representation. In this respect [Inf] dovetails nicely with both Elgin’s account of scientific modelling as providing an *understanding* of nature that does not involve mirroring, and Batterman’s critique of mapping accounts of mathematical representation.

The schema ‘X represents Y’ is then an idealization, and not a benign one for it suppresses the role of the agent in representation. Ronald Giere [2004] suggests the following replacement:

S uses X to represent Y for purposes P.

Different schemas of course highlight different features of representings, but even this suppresses important features of representational practice. One may not, for example, merely use an image to represent a person, but to represent a person *as* something. Margaret Thatcher, to use van Fraassen’s own example, was often represented as a dragon in satirical cartoons. Or, to use a notorious example, the Danish newspaper *Jyllands-Posten* printed a cartoon in which the Islamic prophet Mohammed was represented as a terrorist. More germane to the current discussion, scientists may represent springs *as* simple harmonic oscillators, or use inscriptions of differential equations to represent concrete systems *as* mathematical structures.⁴⁶ For this reason, I prefer van Fraassen’s more illuminating schema:

Z uses X to represent Y as F

A couple of clarificatory remarks are in order. Firstly, by a ‘concrete system’ I mean any system that is not abstract in the sense defined in the first chapter, *viz.* any system that can make a difference to something else. Secondly, representing Y as F does not involve holding that Y *really is* (fundamentally, ontologically) F, or that F *really is* (fundamentally, ontologically) Y. Thus, representing concrete systems as mathematical structures does not, for instance, involve *reinterpreting* mathematics so that it is the study of concrete structures rather than abstract objects, *à la* Hellman. The analogy with pictorial representation is here instructive: representing Winston Churchill as a bulldog, for instance, does

⁴⁶The centrality of representation-*as* in the sciences is a theme emphasized by R.I.G. Hughes.

not require on the part of the representing agent the belief that Churchill really is canine. To give an unadorned example, we will take a concrete system C to be represented by a mathematical structure C^* , where C^* is an n -tuple of real numbers. If, for instance, C is some physical body, $C^* = \langle x_1, x_2, x_3 \rangle$, where x_1 is the mass in kg of C , x_2 is the acceleration in ms^{-2} of C , and x_3 is the force in N on C .⁴⁷ Representing C as the mathematical structure C^* allows it to be brought under the governance of equations, for instance Newton's famous Second Law, $F = ma$. $F = ma$ is itself more clearly parsed as $\forall x[\text{force}_N(x) = \text{mass}_{\text{kg}}(x) \times \text{acceleration}_{\text{ms}^{-2}}(x)]$. These equations allow us to make inferences about C^* . C^* is tethered to the concrete world via measurement procedures; concrete states, processes or events by which one can assign numbers to C , and hence to represent it *as* C^* . In this way, the theory Γ that uses C^* to model C is made subject to empirical confirmation or disconfirmation—confirmation or disconfirmation *as nominalistically adequate*. It is the combination of the practices of measurement and the use of equations that govern the inferences we make about C^* , that allow us to reason surrogatively about C itself. C^* is to C as the bulldog is to Churchill: C^* *itself* plays no role in this practice, only inscriptions of equations that are used to represent it. Nor is C^* *itself* required to explain the *efficacy* of the practice: the practice is efficacious because *thinking about* C^* facilitates surrogative reasoning about C . Once one recognises that representation is not a relationship that obtains between a source and a target or representation, but an act performed by agents, the applicability of mathematics to the sciences is no longer mysterious from a nominalistic point of view. Representing a concrete system as a mathematical structure, requires agents to *represent* mathematical structures, whether those structures *exist* is orthogonal to this practice. Metaphysical theories of representation require there to *be* relevant mappings from the source to the target of representation, but we have seen that the existence of these functions is neither necessary nor sufficient for representation to take place. It may well be the case that in particular instances of representation, agents must themselves be able to represent functions from mathematical structures to concrete systems in order to utilize the derivations they make in the mathematical domain to infer consequences about the concrete system they are representing. But, crucially, this involves mathematical inscriptions, speech acts and thoughts, *not* mathematical objects. We can make vivid this idea by means of an illustration. Suppose that a group make use of a sophisticated clockwork orrery to model the solar system. By observing its movements they can make accurate predictions about the relative locations of the planets and their moons which they then go on to check empirically. One day disaster strikes; the orrery is stolen and melted down for scrap. But one of the group is endowed with a praeternatural memory and knows the working of the device by heart. So precise is this knowledge that she can predict its movements as accurately as if it were there to be observed. In order for the other group members to continue in their representational practice, the mnemonically gifted group member provides a series of descriptions of the no-longer-extant orrery so that her colleagues can continue in their representational practices as she does. The practice of modelling goes on, absent the orrery. What

⁴⁷This is a one-dimensionalist idealization. Forces and accelerations are represented by vectors, themselves represented by triples of reals. Hence $C^* = \langle x, \langle a_1, a_2, a_3 \rangle, \langle f_1, f_2, f_3 \rangle \rangle$ where a_1 – a_3 represent the acceleration of C , and f_1 – f_3 represent the force of C .

was required for these practices was never the orrery itself, only the ability to represent it linguistically, verbally and in thought. Mathematical objects are not required to make sense of our ability to represent the world mathematically, although, as with the orrery, in carrying out these representational practices, they are quantified over.

Philosophers often talk of the ‘role of mathematical objects in science’ or in terms of persons making use of mathematical objects in some way. Baker, in stating the second premiss of his Enhanced Indispensability Argument, talks in this way: ‘Mathematical objects’, he says, ‘play an indispensable explanatory role in science’. Mark Colyvan, in his monograph *The Indispensability of Mathematics* also has a tendency to talk this way.⁴⁸ While it is unlikely that either Baker or Colyvan really believe that people use mathematical objects *themselves*, as though one might use a hammer to drive a nail into a piece of plywood, to talk this way infelicitously makes it appear as though mathematical *objects* are implicated in mathematical *practices* in a way in which they are not. Van Fraassen’s schema ‘Z uses X to represent Y as F’ makes this clear. Scientists use, as we have said, mathematical inscriptions, speech acts and thoughts to represent concrete systems as mathematical objects. Thus far then one need not postulate mathematical objects to make sense of the applicability of mathematics.

A further point stemming from our discussion of scientific representation is that, since scientific representation does not require particular metaphysical relations to hold between the source and the target of representation, one *need not* conclude—as may appear to be the case from Bueno and Colyvan’s inferential account, or Bueno and French’s partial morphism account—that relationships between mathematical structure and concrete systems need to obtain for mathematical representation to take place. Not only is nothing of the sort *required*, nothing of the sort is so much as *relevant* to our ability to represent the world mathematically. At most, what is needed is that agents can represent structural relations between mathematical objects. This much is sufficient to make inferences about concrete systems which can then be tested by means of measurement procedures—all of which is nominalistically acceptable. I have said however that these formal accounts of applied mathematics do get *something* right, but what is it, and what is the relationship between Suárez’s account of scientific representation and the formal accounts of Bueno, Colyvan and French? At a first pass it is this: Suárez’s inferential account tells us what is required for representation to take place, whereas the formal accounts tell us what is required for this representation to be a *good* one. This story however is not quite right, and, for that matter, is committed to the existence of mathematical objects. A better account is this: Suárez’s inferential account tells us what is required for representation to take place, whereas the formal accounts make explicit the structural features that mathematical objects would have *if* they existed, and, by doing so, explain why *reasoning about* those mathematical objects allows for successful surrogative reasoning about concrete systems.

⁴⁸A tendency criticized by Julian Cole and Stewart Shapiro in their review of the book: ‘It is mathematical statements, not mathematical objects, that are (or are not) indispensable to science. The Quinean infers the existence of mathematical objects via ... the univocity of existence. But we can conclude this only if we interpret mathematical language in a certain way—by taking its surface grammatical structure at face value.’ [Cole & Shapiro 2003: 333]

One need not postulate mathematical objects in order to forge a philosophical understanding of mathematical representation in the sciences. The efficacy of these practices makes sense in a universe without abstracta. One final *ad rem* remark is in order. *Mathematical sentences do not function as assumptions in the practices of scientific representation.* They are used of course—used to represent concrete systems as mathematical structures—but assertion is not the sole use to which sentences can be put. Users of mathematical sentences must gasp their content in order to make inferences with them, but belief that they accurately describe a domain of abstract objects is optional. The role of mathematical sentences in our cognitive ecology is more like that of make-believe.⁴⁹ A defect then of theories of scientific explanation such as the deductive-nomological account is that, by treating every sentence used by scientists an assumption, they cannot account for this aspect of the application of mathematics. That is, if one treats all the sentences we make use of in a theory as asserted as true (as opposed to merely nominally adequate), one cannot account for representing something *as* something else; for representing something as something else does not require *beliefs* regarding the properties of that something else, only about the properties it would have were it to exist. Putnam could not be accused of an unsophisticated understanding of science or its philosophy, but his no-miracle argument glosses over the details of scientific practice that would allow one to assess its cogency. The consequences of a proper understanding of scientific representation are utterly deflationary, for not only can one understand the applicability of mathematics without postulating mathematical objects, those who do engage in applied mathematics can deny the existence of mathematical objects without engaging in intellectual doublethink.

Earlier we criticized the Quinian method of ontology on the grounds that its adoption entailed a commitment to a strong form of ontological anti-realism. I want to add to this the criticism that the Quinian method of ontology *mislocates* where ontological commitments are to be found. The devotee to the Quinian method will not be impressed by the preceding discussion. She will claim that if mathematical explanations involve comparisons between mathematical structures and physical systems, then our philosophical account of mathematical explanations still necessarily involves quantification over mathematical objects, and so a commitment to their existence. Some authors attempt to sidestep quantification over mathematical objects by appealing only to *properties* of those objects. Roman Frigg [2010] endorses this approach in his defence of a fictional account of abstract models in the sciences. His concern is with how to understand ‘transfictional’ statements—statements that make comparisons between putatively fictional and real objects:

[T]ransfictional statements about models should be read as prefixed with a clause stating what the relevant respects of the comparison are, and this allows us to rephrase comparative

⁴⁹This is put well by Sygne and Griffith, in a textbook on classical mechanics:

In fact, mechanics—and indeed all theoretical science—is a game of mathematical makebelieve. We say: If the Earth were a homogeneous rigid ellipsoid acted on by such and such forces, how would it behave? Working out the answer to this mathematical question, we compare our results with observation. If there is agreement, we say we have chosen a good model; if disagreement, then the models or laws assumed are bad. [Synge & Griffith 1959: 5]

sentences as comparisons between properties rather than objects, which makes the original puzzle go away. Hence, truth conditions for transfictional statements (in the context of scientific modelling) come down to truth conditions for comparative statements between properties ... For instance, when I say ‘my friend Peter is just like Zapp’ I am not comparing my friend to a nonexistent person. What I am asserting is that both my Peter and Zapp possess certain relevant properties (Zapp possesses properties in the sense explained above) and that these properties are similar in relevant ways. [Frigg 2010: 263]

Comparisons between fictional and real objects are thought to be problematic because their truth conditions would require domains which include fictional objects. Frigg avoids this by holding that all that is required is a domain of properties; the similarity claim is *true* so long as there are certain properties that are instantiated in both Peter and Zapp. My claim is that seeking truth conditions for such statements is an unavailing task. Scientists represent the world mathematically; laypeople represent the world using fictions. As philosophers our goal is to understand the preconditions for these practices; what must one suppose is the case to make sense of how these things are possible? Positing objects to satisfy that which people quantify over does not achieve this; as we have seen in the case of representation. To ask what is quantified over is to ask the wrong question. A better question is this: *what must exist in order for successful mathematical representation of the world to take place?*⁵⁰

5.3.6 How Modelling Provides Epistemic Access to the World

The mere applicability of mathematics is not mysterious to the nominalist. Whether or not mathematical objects exist, I have argued, is simply beside the point in a sound account of applied mathematics. I want to say something now about why this process affords *knowledge*; how mathematical modelling provides epistemic access to the world. The manner in which modelling provides epistemic access to the world can be embedded in a broader story—put forward by Brandom [1994, 2008]—about the manner in which interacting with the world can act to fashion and refashion our commitments about it. Central to this story are trial and error procedures, characterized by ‘the cycle of perception, inference, action and perception’ [Brandom 1994: 332]. In Brandom’s scenario, a group of language users employ the term ‘acid*’. The circumstances of application for acid* are when an item tastes sour, and its consequences of application are that an acid* will turn phenolphthalein blue—i.e. if an item tastes sour, these language users are entitled and committed to applying the term acid* to it, and given any item that has been termed an acid*, these language users are entitled and committed to its turning phenolphthalein blue. Acid* then is characterized by two conditionals:

(circ): If an item is sour then it is an acid*.

(cons): If an item is an acid* then it turns phenolphthalein blue.

⁵⁰This is a thought that will be made more precise later in this chapter.

In employing the term ‘acid*’, this community of language users are implicitly endorsing these inferential commitments. If a member of the community finds an item that tastes sour but which turns phenolphthalein red, then she is faced with mutually incompatible commitments: she infers from the items sour taste that it is an acid*, and from its being an acid* to its propensity to turn phenolphthalein blue, yet by exercising her capacities of perception she finds that the phenolphthalein is red. The cycle of perception, inference, action and perception has resulted in this incompatibility. Faced with the incompatibility, the agent is required to alter at least one of her commitments; she must—the force of this ‘must’ is normative—either give up the belief that the item is sour, the belief that the phenolphthalein is red, or else amend her concept of acidity* to accommodate her perceptions. This may take the form of restricting the circumstances of application, perhaps to *liquids* that are sour or liquids of a certain colour, or the consequences of application, perhaps to turning phenolphthalein blue only under particular circumstances, or even rejecting the concept acid* altogether as defective. In this manner our interactions with the world are able to structure the content of our beliefs. When this process goes aright, the normative structure of one’s beliefs—the network of commitments and entitlements, what one is entitled or committed to infer from what—tracks the modal or nomic structure of the world. This is not to say that our minds become a *mirror* of nature, but rather that they embody an *understanding* of the world by tracking (some of) its modal or nomic properties.

Through this process, the agent in Brandom’s toy example is, to use van Fraassen’s terminology, assigning objects a position in *logical space* in the most restricted sense of this phrase. She is attributing to them a particular inferential structure, one which can be characterized by means of conditionals. Mathematical representation is an elaboration of this process. By representing the world mathematically, one can make more fine-grained the normative inferential structure of one’s commitments than would be practically possible if that structure consisted only in conditionals; the sorts of logical spaces that are available to agents is increased. For instance, by making use of differential equations—Newton’s second law, say—one can adopt more sophisticated inferential commitments. Commitments of this sort do not have a finite set of circumstances or consequences of application; instead measurement procedures take the place of straightforward perceptions, and allow values to be assigned to F , m and a . The values these take will then determine the consequences of application, and these consequences can in turn be perceived. This may result in an incompatibility within an agent or group of agents’ commitments. In which case at least one of these commitments must be revised, and the process starts anew. Empirically testing scientific hypotheses is a highly sophisticated cycle of this form.⁵¹

⁵¹It may be thought that this fails to provide an account of how mathematics can be used *representationally*, as opposed to merely make inferences about concrete systems. I am enough of an inferentialist to see no distinction between the two. To represent something as thus and so *just is* to endorse certain inferences regarding it. To take some object o to be red, for instance, is to endorse certain inferences about it; that o is not green, that o is coloured, and so forth. Conversely, to endorse inferences about something, just is to characterize that thing as thus and so. Inferring and representing are two sides of the same intentional coin.

Metaphor and Thought

One can elaborate this account by examining the role of metaphorical or analogical thinking—an ability to abstract relational patterns from individual instances and to recognize further instances of these patterns; i.e. to *apply* them in novel contexts—in human⁵² cognition. The classical theory of metaphor, and the notion of metaphor that most laypeople take for granted, holds metaphors to be exceptional occurrences in the language; special devices employed by poets or the poetically inclined which generally occur only in literary contexts. At variance with this picture however is the theory of metaphor accepted by contemporary linguists. Brian Bowdle and Dedre Gentner state, with regards to the theory of metaphor:

Over the past two decades, the cognitive perspective on metaphor has undergone a radical shift. Traditionally, metaphors have been treated as both rare in comparison to literal language and largely ornamental in nature. Current research suggests precisely the opposite. Rather than being restricted to poetic uses, metaphor is common in everyday communication. [Bowdle and Gentner 2005]

They continue:

A growing body of linguistic evidence further suggests that metaphors are important for communicating about, and perhaps even reasoning with, abstract concepts such as time and emotion. [...] Indeed, studies of scientific writing support the notion that far from being mere rhetorical flourishes, metaphors are often used to invent, organize, and illuminate theoretical constructs. [ibid.]

Recent statistical research by linguists has demonstrated that much of our conventional language is in fact metaphorical in nature, usurping the assumption that ordinary language was most often literal, and that metaphor was only employed peripherally. George Lakoff has collated a list of such conventional metaphors which extends *ad nauseam*,⁵³ numerous as the examples are. But we will limit our attention to one: LOVE AS A JOURNEY. English is littered with phrases that employ this metaphor⁵⁴:

- Look how far we've come.
- It's been a long, bumpy road.
- We can't turn back now.
- We're at a crossroads.
- We may have to go our separate ways.
- The relationship isn't going anywhere.
- We're spinning our wheels.

⁵²Whilst analogical thinking develops spontaneously in human infants, chimpanzees require training to think analogically, and lesser apes such as monkeys are seemingly unable to think analogically [Thompson and Oden 1998].

⁵³C.f. <http://cogsci.berkeley.edu/lakoff/MetaphorHome.html>

⁵⁴I take the examples from Lakoff [1992].

- Our relationship is off the track.
- The marriage is on the rocks.
- We may have to bail out of this relationship.

Such use of metaphor permeates our ordinary speech; in the previous paragraph alone the words *usurping*, *employed*, *peripherally*, *extends* and *littered* were all used in a metaphorical sense. Not only does the contemporary theory of metaphor diverge from the classical theory in what it asserts about the prevalence of metaphor in ordinary language, it differs in what it takes metaphors to *be*. An analogy or metaphor (a subspecies of analogy), is not a word or expression, it is a mapping from one conceptual domain to another. As Gentner states:

The analogy “A *T* is (like) a *B*” defines a mapping from *B* to *T*. *T* will be called the *target*, since it is the domain being explicated. *B* will be called the *base*, since it is the domain that serves as a source of knowledge. [Gentner 1983: 157]

Analogy and metaphor, as opposed to simple similarity comparisons, tend to preserve, in their mappings, relations between objects (i.e. mappings between dyadic or polyadic predicates), whilst discarding attributes of objects (i.e. mappings between monadic predicates):

Disregard properties $A(b_i) \rightarrow A(t_i)$

Preserve relations $R(b_i, b_j) \rightarrow R(t_i, t_j)$

In the LOVE AS A JOURNEY metaphor, the source domain is JOURNEYS and the target domain is LOVE. Clearly, mappings from monadic predicates are disregarded: by saying ‘It’s been a long, bumpy road’ one does not intend to communicate that their relationship is made from pothole-riddled asphalt. Rather, the relation between travel and long, bumpy roads is preserved; i.e. that long, bumpy roads make travel difficult.

A metaphor then is not merely an article of language, but a way in which we structure our thoughts and reasoning processes. Metaphor can also be used as an inferential aid; for instance, classical categories can be understood in light of the metaphor of physical containers. We often talk of things being in a certain category (Battleship Potemkin is ‘in’ the category of revolutionary propaganda films) and of things being removed from certain categories (Robert Mugabe has been ‘removed’ from the category of Edinburgh alumni). How then does this aid inferences? Lakoff explains:

Under the CLASSICAL CATEGORIES ARE CONTAINERS metaphor, the logical properties of categories are inherited from the logical properties of containers. [...] Thus, the logical properties of classical categories can be seen as following from the topological properties of containers plus the metaphorical mapping from containers to categories. As long as the topological properties of containers are preserved by the mapping, this result will be true. [ibid., p14]

By conceptualizing physical containers, we can make inferences about classical categories. This can be seen with the example of the classical syllogism: The Creature is an oyster. All oysters are bivalve molluscs. Therefore, the creature is a bivalve mollusc. The same inference could be made by visualizing the topological properties of physical containers: The Creature is in the *oyster* box, and the *oyster* box is inside the *bivalve mollusc* box, therefore The Creature is inside the *bivalve mollusc* box.

The close parallel between the contemporary theory of metaphor and both Suárez's inferential theory of scientific representation and Bueno and Colyvan's inferential conception of applied mathematics—in which 'the idea ... is that the mapping carries the inferential structure from one domain to the other—should be clear.⁵⁵ Just as in the *immersion* stage of mathematical applications, a mapping is made from the source domain under investigation to a target domain which we use to comprehend the former. As in the *derivation* and *interpretation* stages, consequences are drawn in the target domain which are then interpreted as results in the source domain. Metaphors, like mathematical structures, can come in varying degrees of aptness for the source domain they are used to explicate. George Bush, for instance, was criticised for introducing the metaphorical term 'axis of evil' to describe the states of Iran, Iraq and North Korea on the grounds that it is not apt (read: the topological structure of an axis is not homomorphic to the inferential structure of the evil associated with Iran, Iraq and North Korea circa 2002). Similarly, a Euclidean theory of space, for instance, could be criticised on the grounds that it is not apt insofar as the Euclidean structure employed by the theory is not—or, better, *would not be, were it to exist*—homomorphic to our non-Euclidean space.

What lessons are we to draw from contemporary theory of metaphor? Firstly, the great pervasiveness of metaphorical modes of thinking is itself significant⁵⁶. If much of our conceptual system is indeed analogical in nature, then we could expect science, being a sophisticated part of that conceptual scheme, to exploit metaphorical techniques. Secondly, there is nothing mysterious about how mathematical thinking provides epistemic access to the world. Analogical thinking is an uncontroversial part of our cognitive apparatus, and is studied *naturalistically* by cognitive scientists. No appeal to non-natural process or entities, including of course mathematical objects, is required to make sense of analogical thinking or the knowledge it affords us. Mathematical modelling provides epistemic access to the world in the same way that metaphors do. Metaphors provide us with an inferential grasp of something in virtue of our inferential grasp of something else. Metaphors and models allow us to understand new domains by allowing us to *reason* about them. In the scientific case, whether the metaphor is a felicitous one is tested empirically through measurement procedures. A puzzle generated by nominalism was how

⁵⁵Recall Suárez's [inf] condition:

[Inf] *X* represents *Y* only if (i) the representational force of *X* points towards *Y*, and (ii) *X* allows competent and informed agents to draw specific inferences regarding *Y*. [Suárez 2004: 773]

⁵⁶There is a limit however to the extent to which metaphorical thinking structures and prescribes our thought. Lakoff and Johnson [1980b and 2000] utilise this theory of metaphor for rather extravagant ends, arguing, for instance, that there is no objective truth, merely competing metaphors. We need not follow Lakoff and Johnson to these extremes; all that we require is the fact that people create conceptual mappings between domains in order to understand and articulate facts about one domain in terms of another. Such processes can be prevalent in our thinking (and the linguistic evidence appears to confirm this) without *dictating* our thinking. See Steven Pinker [2007] for a nice summary of these issues.

it could be that something false could tell us anything about the world, but if mathematical modelling is akin to metaphorical grasping as I have argued it is, then the puzzle is resolved. What is required for mathematics to provide understanding about the physical world is not mathematical *knowledge*, but mathematical *understanding*.

5.4 What is Explanation?

There is nothing involved in a philosophical understanding of scientific representation that requires the existence of mathematical objects. But perhaps there is something special about cases in which applied mathematics is *explanatory*, as Baker suggests. Recall that Baker motivates this notion through a comparison of mathematical explanation with causal explanation. The comparison, as we noted, is illicit: mathematical explanation does not, in all respects, resemble causal explanation. If it did, the existence of mathematical objects would not be contested. If we are to ascertain whether there are genuinely explanatory applications of mathematics in the empirical sciences then what is required is an account of *explanatoriness*. The debate on the role of mathematical explanations in science has not included any serious attempt to carry out this task, but without some exposition of the nature of explanation there is no principled way in which one can hope to settle the two questions in this area that are relevant to our ontological dispute: whether there are in fact authentic mathematical explanations of concrete phenomena, and whether their explanatoriness provides a special reason to think that the mathematical objects quantified over in the explanations actually exist. What is required then are some minimal conditions for explanatoriness.

There are, I suggest, at least three conditions of adequacy on any such minimal account of explanation. Firstly, it must account for the sheer diversity of explanation. As was noted earlier, explanations come in myriad flora and fauna; at the very least *Causal Explanation*; *Theoretical Reduction*; *Conceptual Explanation*; *Unification*; *Structural Explanation*; *Deductive-Nomological Explanation*; *Teleological Explanation*; *Statistical/Probabilistic Explanation*. Any minimalist account of explanation must account for this diversity; it cannot rule out any of these types of explanation, whilst, at the same time, uncovering what is common to them all. Secondly, it must make sense of the epistemological character of explanation; that is to say it must account for the differences that explanations can make to the epistemological predicament of agents. It must show how explanation is related to understanding. Thirdly, the account must make sense of the *illocutionary* nature of explanation. This is a point made by Peter Achinstein [1983]. Imagine for example that a person S explains an event by citing a cause of that event. S utters the sentence:

(*) George has high levels of serotonin in his nervous system because he has been administered a selective serotonin re-uptake inhibitor.

The utterance of (*) in this context is indeed an explanation, but to make an utterance that expresses a

proposition which cites the cause of an event is not necessarily to provide an explanation of that event. Whether such an utterance is tantamount to an explanation depends on the intentions of the speaker. An utterance of (*) for instance in response to the question

Who has high levels of serotonin in their nervous system in virtue of being administered a selective serotonin re-uptake inhibitor?

would not constitute an explanation, but rather an identification of the person whose serotonin levels are elevated. Causal propositions are not then “intrinsic explainers” whose expression is alone sufficient to constitute an explanation.⁵⁷ (*) has the status of an explanation in the first context because it is uttered with the intention to produce *understanding* in the audience who hear it, with respect to the serotonin levels of the person in question. The point is quite general; sentences, sets of sentences or ordered n-tuples of sentences are not *inherently* explanations.

5.4.1 Prominent Models of Explanation

It might be hoped that one or more of the prominent models of explanation in the literature will provide the minimal conditions for explanation that we require. Causal accounts—such as Peter Lipton’s [2004], Wesley Salmon’s Causal Mechanical (C-M) model [Salmon 1984, 1994, 1997], or Baruch Brody’s extension of Hempel’s Deductive-Nomological (D-N) model [Hempel 1965] with a causal condition [Brody 1972]—are clearly out of the question. This is not because, as Baker [2005] suggests, that to adopt a causal account of explanation would necessarily beg the question against the platonist. If a merely causal account of explanation can be independently motivated then its adoption would be perfectly legitimate, regardless of whether this undermined one putative argument for platonism. The reason that causal theories of explanation will not provide the minimal conditions we are seeking is because there are explanations that are not causal. Lipton provides a neat example:

[S]uppose that a bunch of sticks are thrown into the air with a lot of spin, so that they separate and tumble about as they fall. Now freeze the scene at a moment during the sticks’ descent. Why are appreciably more of them near the horizontal axis than near the vertical, rather than more or less equal numbers near each orientation as one might have expected? The answer, roughly speaking, is that there are many more ways for a stick to be near the horizontal than near the vertical. ... There are infinitely many horizontal orientations, but only two vertical orientations. ... Another way of putting it is that ... there are two horizontal dimensions but only one vertical one. [Lipton 2004: 31-2]

Causal accounts of explanation must be set to one side. More promising for our purposes then is the D-N model, absent any additional causal condition. According to the D-N model explanations are *arguments* whose premises include nomic laws and whose conclusion is the explanandum. For instance:

The phenolphthalein was exposed to acid.

⁵⁷This is not to claim that an audience with the right sort of interests and abilities could not *make use* of (*) to explain the high levels of serotonin in the subject, or for that matter other explananda.

All phenolphthalein when exposed to acid becomes colourless.

∴ The phenolphthalein became colourless.

This is a deductively valid argument and, provided the premisses are true, looks to be an authentic explanation. Unfortunately the D-N model is both too strong and too weak. It is too strong because many explanations do not take the form of a deductive argument with nomic laws in the premisses, particularly in the special sciences and in everyday contexts, as well as in mathematics. Consider straightforward causal explanations such as:

(Brick) The window smashed because a brick was thrown at it.

(Brick) is an explanation but is not an argument of any form, let alone one with nomic laws in its premisses. One response (Hempel's response in fact) is to hold that *superficially* non-D-N explanations are in fact *implicit* D-N explanations, and can be appropriately translated. This is implausible at best when taken to apply to conceptual explanations, such as a conceptual analysis of knowledge, but even if all explanations could be so transformed there remains a problem. Assume with Hempel that the "real" explanation underlying (Brick) is the following (Brick*):

Whenever a brick is thrown at a window, and further conditions Θ are met, the window will smash.

A brick was thrown at the window and further conditions Θ were met.

∴ The window smashed.

Our second condition of adequacy is that a minimal account of explanation must accommodate the *epistemological* character of explanations; their connexion to understanding. If *prima facie* explanations such as (Brick) are not really explanations, then they would not furnish agents with understanding; only real explanations could provide this expository cognitive role. Agents however are not generally *cognizant* of the D-N counterparts to the explanations they actually make use of—even in the simple case of (Brick*), few could adumbrate the further conditions Θ required to grasp the "real" explanation—so the proposed solution is incompatible with the fact that our ordinary explanatory practices produce understanding. The D-N model is too weak because it allows for trivial explanations. One could provide a sound deductive-nomological argument that the phenolphthalein turned colourless by including the conclusion in the premisses along with *any* nomic premiss, Coulomb's law perhaps. The model is also too *weak* because even non-trivial, deductively sound nomic arguments are not necessarily explanatory. As Lipton notes: 'We may explain why an object was warmed by pointing out that it was in the sun and everything warms when in the sun, but we cannot explain why an object was not in the sun by pointing out that it was not warmed' [Lipton 2004: 27]. Explanations are not (except contingently) arguments; the necessary and sufficient conditions for good arguments are not the necessary and sufficient conditions for good explanations.

A second avenue of possibility is Salmon's Statistical-Relevance (S-R) model.⁵⁸ The guiding idea here is that statistically relevant relationships between classes of events are explanatorily relevant. Scientific explanations in general are, on this picture, explanations of why an event X , which is a member of some class A , also a member of class B . An explanation then is a set of empirical probabilities relating A and B and a statement that the explanandum X is a member of a relevant class. This will take the following form:

$$\begin{aligned} P(B|A \cap C_1) &= p_1 \\ P(B|A \cap C_2) &= p_2 \\ &\cdot \\ &\cdot \\ &\cdot \\ P(B|A \cap C_n) &= p_n \\ X \in C_k (1 \leq k \leq n) \end{aligned}$$

Here, $A \cap C_1, \dots, A \cap C_n$ are *partitions* of A , meaning that $\bigcup_{i=1}^n A \cap C_i = A$ and $\bigcap_{i=1}^n A \cap C_i = \emptyset$; *viz.* the partitions are exclusive and exhaustive. Two further conditions must be met. The probabilities p_1, \dots, p_n must all be different, and $A \cap C_1, \dots, A \cap C_n$ must be *homogeneous* with respect to A . A homogeneous partition is one in which for all C_i, C_j such that $C_i \neq C_j$, $P(B|C_i) \neq P(B|C_j)$, and there is no additional statistically relevant partition of A , i.e. there is no additional set D such that for some C_i $P(B|C_i) \neq P(B|C_j \cap D)$. Returning to the question of why the phenolphthalein turned colourless, we can construct an explanation using the S-R model:

$$\begin{aligned} A &= \text{the class of chemical compounds} \\ B &= \text{the class of things that turn colourless when exposed to acid} \\ C_1 &= \text{the class of things that are phenolphthalein} \\ C_2 &= \text{the class of things that are not phenolphthalein} \end{aligned}$$

$$\begin{aligned} P(B|A \cap C_1) &= 1 \\ P(B|A \cap C_2) &= 0 \\ X \in C_1 \end{aligned}$$

As above, this looks to be a genuine explanation. One concern we might have here is that statistical relevance is intended to capture *causal* relevance. This is why, facing objection that it could not do so, Salmon went on to abandon the S-R model in favour of the explicitly causal C-M model. If this is the case then the S-R model will too lack the universality we require. A more pressing problem is that it fails to capture the epistemological character of explanation for the same reasons as the D-N model. Not all explanations take this form, but even if, contrary to appearances, every explanation *could* be composed so as to do so, people in general would not be aware of the recomposed counterparts to the explanations they actually make use of; in which case the link between explanation and understanding is again severed.

⁵⁸I am indebted to Achinstein's [1983] exposition.

It might finally be hoped that *unificationist* models of explanation, such as those proposed by Michael Friedman [1974] and Philip Kitcher [1989] will provide some minimal account of what is required for explanatoriness. In Kitcher's model, explanation is the process of describing as much of our knowledge as possible using as few argument patterns as possible. Kitcher precisifies this notion as follows. A *schematic sentence* is sentence scheme that is obtained by replacing some of the nonlogical vocabulary with variables. So, for instance, a schematic counterpart to 'All phenolphthalein turns colourless when exposed to acid' might be 'All phenolphthalein turns Y when exposed to acid' or 'All X turns Y when exposed to Z'. This is complemented with a set of *filling instructions*, specifying how variables are to be filled in with ordinary expressions. A *schematic argument* is a sequence of sentence schemes and an argument *classification* indicates which terms in a schematic argument are premises, and which sentences can be inferred from which others, given which rules of inference. A *general argument pattern* is an ordered triple which includes a schematic argument, a set of filling instructions and a classification for the argument. Let K be the set of accepted scientific beliefs (idealized to be consistent). Then $E(K)$ is the *explanatory store*: the set of argument patterns that best unifies K ; i.e. describes as broad a range of phenomena as possible using as few argument patterns as possible. Unification looks more promising as an account of theoretical reduction and conceptual explanations than either the D-N or the S-R models. Unfortunately, this kind of unificatory theory is subject to the same epistemological objection that afflicts the D-N and S-R models. Unification, of the sort proposed by Kitcher, cannot provide necessary conditions for explanation because it is too strong; there are explanations, such as individual causal or mathematical explanations, that do not result in any unification in Kitcher's sense. Even if counterparts to these explanations can in fact be captured by the unificatory model, it is psychologically implausible that agents arrive at these causal explanations by comparing explanatory stores to find that which best systematizes K .

While all of the examined theories of explanation may pick out interesting subclasses of our explanatory practices, all are too strong (and for that matter too weak) to act as minimal conditions for explanatoriness. As such, they are not of use in assessing whether there are mathematical explanations of empirical phenomena. A better route, I claim, is to approach explanation through the cognate notion of *understanding*, and it is to this which we shall now turn.

5.4.2 Explanation and Understanding

How are we to expound understanding? There is a burgeoning body of work amongst epistemologists attempting to make sense of understanding. Whilst their primary goal has not been to say what understanding *is*, but rather to settle questions regarding whether understanding is a species of knowledge and why understanding is valuable, their intuitive characterizations of understanding are suggestive, and can serve as a starting point for our discussion of this more foundational question. Jonathan Kvanvig suggests the following:

The central feature of understanding, it seems to me, is in the neighborhood of what internalist coherence theories say about justification. Understanding requires the grasping of explanatory and other coherence-making relationships in a large and comprehensive body of information. One can know many unrelated pieces of information, but understanding is achieved only when informational items are pieced together by the subject in question. [Kvanvig 2003: 192]

And:

[U]nderstanding requires, and knowledge does not, an internal grasping or appreciation of how the various elements in a body of information are related to each other in terms of explanatory, logical, probabilistic, and other kinds of relations that coherentists have thought constitutive of justification. [ibid.: 192-3]

In a similar vein, Catherine Elgin states:

Understanding, then, is in the first instance a cognitive relation to comprehensive, coherent sets of cognitive commitments. The understanding individual propositions express derives from an understanding of larger bodies of information they belong to. [...] The understander must also grasp how the various truths relate to each other. This is an important point. One might think that the comprehensive body of information is just a large collection of propositions. I suggest that understanding involves more. The understander should be able (and perhaps be aware that she is able) to use that information—for example, to reason with it, to apply it, to perhaps use it as a source of working hypotheses about other related matters. [Elgin 2009: 323]

Clearly, if we are to define explanation *in terms of* understanding in a useful way then we cannot adopt Kvanvig's explication wholesale, lest our explicandum include its explicans. However, something of the sort adumbrated by Kvanvig and Elgin is surely on the right tracks.

Semantic Inferentialism and Understanding

In fact, a fuller conception of understanding along the lines of Kvanvig and Elgin's glosses has been worked out, although not in a (primarily) epistemological context. For this, we must turn instead to the philosophy of language. Robert Brandom, in a number of places, but primarily in his tome-like opus *Making it Explicit* and two shorter monographs *Articulating Reasons* and *Reason in Philosophy*, has expounded a view of understanding that acts as the lynchpin of his theory of linguistic meaning, *semantic inferentialism*. My goal here is not to defend semantic inferentialism, although I find the doctrine attractive, but rather to cash out the purely epistemological sense of understanding intimated in the passages of Kvanvig and Elgin.⁵⁹ Brandom's concern is with how beliefs come to possess their particular content; what it is in virtue of that a given belief means one thing rather than another. To take a contrasting example, what one might call *robustly empiricist* theories of mental content, such as

⁵⁹It is however a happy consequence of adopting semantic inferentialism, on my view, that doing so admits of a *uniform* account of understanding language and understanding more generally.

Dretske's *indicator semantics*⁶⁰ and Millikan's *teleosemantics*⁶¹ are attempts to explain the content of a given belief in terms of its causal aetiology. These theories all essentially involve an ability to reliably differentially respond to events in the environment. Animals that make reliable differential responses to events of certain sorts have neuro-physiological structures that covary appropriately with particular event types. The condition is supposed to secure the link between the content of a belief and its being true since, on this account, what makes an neuro-physiological event about, say, red objects is that it covaries with the presence of red objects; it carries the information that red objects are present. This is a necessary but not sufficient condition for a creature to possess content, for although we have reliable differential responses to, say, heat, and we have beliefs about heat, thermometers reliably differentially respond to heat but lack beliefs about it. So too do lumps of wax reliably differentially respond to heat by melting, chunks of iron to rain by rusting, and so on. As such, sophisticated empiricists build in other prerequisites in order to specify sufficient conditions for content.⁶² Reliable differential responses are not sufficient for belief however: a trained parrot's reliable disposition to respond differentially to red stimuli by squawking "Red!" is not sufficient for it to have a true belief about the colour of the objects it is reliably responding to nor indeed for it to have a belief at all—even if it superficially resembles believing creatures in some respects. What missing ingredient then is required to elevate the creature that reliably differentially responds to various stimuli to a creature that has *beliefs regarding* the stimuli to which it responds? One answer begins by noting that the parrot lacks *comprehension*: it has no mastery, no understanding of redness; it lacks the concept RED. This is alright as far as it goes, but we have so far only succeeded in giving a series of labels to the parrot's deficiencies. If we are to understand what it is lacking, we must somehow unpack the consanguineous notions of *comprehension*, *understanding* and *concept possession*. Semantic inferentialism is, at its core, an attempt to do just this. Robert Brandom provides the canonical statement of the view:

To grasp or understand a concept is ... to have practical mastery over the *inferences* it is involved in—to know, in the practical sense of being able to distinguish, what follows from the applicability of a concept and what it follows from. The parrot does not treat "That's red" as incompatible with "That's green," nor as following from "That's scarlet" and entailing "That's colored." Insofar as the repeatable response is not, for the parrot, caught up in practical proprieties of inference and justification, and so of the making of further judgements, it is not a conceptual or cognitive matter at all. What the parrot and measuring instrument lack is an appreciation of the significance their response has as a reason for making further claims and acquiring further beliefs, its role in justifying some further attitudes and performances and ruling out others. Concepts are essentially inferentially articulated. Grasping them in practice is knowing one's way around the proprieties of inference and incompatibility they are caught up in. What makes a classification deserve to be called *conceptual* classification is its *inferential* role. [Brandom 1994: 89]

⁶⁰Cf. esp. Dretske [1988, 1995].

⁶¹Cf. esp. Millikan [1984, 1989].

⁶²Theories of content that link the content of a belief to the stimuli to which it reliably differentially responds to in this strong way have problems handling necessary truths. Necessary truths are true in all times and places. As a result neuro-physiological structures cannot reliably differentially respond to the 'presence' of these truths in any interesting way. This has the disquieting consequence that no belief can have the content that, say, all triangles are trilaterals.

Whereas broadly Tarskian⁶³ approaches to semantics give explanatory privilege to *truth conditions*, take comprehension to consist in ‘grasping truth conditions’ (whatever that opaque phrase amounts to) and go on to explain inference in terms of truth conditions, semantic inferentialism begins with a basic notion of inference and uses it to do its explanatory work. Some of the distinctive features of semantic inferentialism can be seen in how it contrasts with robustly empiricist theories. For one thing, whilst it is coherent according to these theories for a creature to have *one* concept in virtue of reliably differentially responding in the right sort of way to only *one* thing; according to the semantic inferentialist picture, this is a radical mistake. Grasping a concept requires a practical mastery of its inferential relations to other concepts—what it follows from, what follows from it, what it precludes and what precludes it. To grasp one concept, a creature must grasp *many*. As Brandom puts it, ‘Cognitively, the grasp of one concept is the sound of one hand clapping’ [Brandom 2000: 49].

Brandom, following Wilfred Sellars’ usage, calls the kinds of inferences that make up the content of a concept *material* inferences. Material inferences are those inferences that are licensed by virtue of the *content* of the concepts involved in the inference. From “Glasgow is west of Edinburgh” one can infer “Edinburgh is east of Glasgow”, from “The paint is red (and monochromatic)” that “The paint is not green”, or from “It is raining outside” that “The ground will be wet”. Grasping a concept involves a practical mastery of the material inferences it is implicated in, an ability to *use* the concept appropriately in the practice of giving and asking for reasons and discerning one’s commitments. Inferences of this sort are not logical inferences, which is to say that the practice of making inferences is not limited to those beings who have some logical mastery. Restricting the domain of inference-making to logical discourses involves treating the above inferences as *enthymematic*; invalid until supplemented with appropriate conditionals. “Edinburgh is east of Glasgow” can only be validly inferred from “Glasgow is west of Edinburgh” with the additional premiss “For all x, y [if x is east of y then y is west of x]. Similarly, the conditionals “For all x [if x is red (and monochromatic) then x is not green]” and “If it is raining then the ground will be wet” are required to make the second and third inferences formally valid respectively. For the semantic inferentialist, material inference is fundamental, logical vocabulary (paradigmatically the conditional) is a means of *making explicit* the inferential relations that we endorse in practice (and, if one should so desire, inferential relations that we do not endorse). Explicit logical competence is not a prerequisite for inferential competence; there can be discursive practices that are not logical practices, but logical vocabulary is a means by which we can codify our discursive practices. By making explicit the inferences we could otherwise only adhere to in practice, logical vocabulary allows us to lay bare our inferential commitments, bringing them under rational scrutiny in a way that would not be otherwise possible. Material inferences are non-monotonic; supplementary or collateral premisses can turn good material inferences into bad ones. ‘It is raining outside, therefore the ground will be wet’ is a good

⁶³I call these ‘Tarskian’ approaches only to make reference to their appropriation of the formal definition of truth in a model, which Tarski bestowed us. Tarski himself did not seek to ground a theory of meaning or understanding using this apparatus: ‘[Tarski’s] goal was *not* to use our antecedently understood notion of truth to *endow* the sentences of previously uninterpreted systems with truth conditions, and hence meaning. It was to *define* restricted truth predicates for already meaningful, and antecedently understood, formal languages.’ [Soames 2010: 38-9]

inference, but ‘It is raining outside *and* the ground is covered with tarpaulin, therefore the ground will be wet’ is a bad one. ‘It is raining outside *and* the ground is covered with tarpaulin *and* the tarpaulin has been perforated, therefore the ground will be wet’ is, on the other hand, a good inference.

The content of a belief is constituted by the material relations it stands in to other beliefs. These relations can be classified into three fundamental kinds: *commitment-preservation*; *entitlement-preservation* and *incompatibility*. A proposition p preserves commitment to a proposition q, if commitment to p requires commitment to q. The proposition ‘That is a vixen’ preserves commitment to ‘That is a fox’, for anyone who is committed to the first is thereby committed to the second. A proposition p preserves entitlement to a proposition q if commitment and entitlement to p provides *prima facie* entitlement to q. A person committed and entitled to the proposition ‘It is raining outside’ is thereby *prima facie* entitled to the proposition ‘The ground will be wet’. Finally, a proposition p is incompatible with a proposition q just in case, if a person is entitled to p then she is not entitled to q. A commitment to ‘The paint is red (and monochromatic)’ precludes commitment to ‘The paint is green’.

Semantic inferentialism involves a broad conception of what can play a role in material inference. The ‘inferential’ connexions relevant to the content of a concept involve not only connexions between separate propositions but also between their *circumstances* and *consequences* of application:

[T]he broad conception includes the possibility of noninferential circumstances and consequences of application. In this way ... the specifically empirical conceptual content that concepts exhibit in virtue of their connection to language entries in perception and the specifically practical conceptual content that concepts exhibit in virtue of their connection to language exists in action are incorporated into the inferentialist picture. ... Conceiving such inferences broadly means conceiving them as involving those circumstances and consequences, as well as the connection between them. [Brandom 1994: 131]

Inferential connexions in this broad sense then include not just connexions between propositions but also afferent connexions—the input of stimuli, reliable differential responses to kinds of objects—and efferent connexions—the agent’s actions. Part of what constitutes our concept of redness is red experiences (the other part is how these red experiences are integrated into a network of relations of commitment-preservation, entitlement-preservation and incompatibility). Brandom calls the view that this broad inferential articulation is sufficient to determine conceptual content ‘strong inferentialism’. This is the view developed and defended by Brandom. ‘Hyperinferentialism’ holds that the inferential connexions that obtain between propositions are sufficient to determine content. Finally, ‘weak inferentialism’ is the view that inferential relations are a *necessary* constituent of conceptual content, but not sufficient. On weak inferentialism, if two contents were involved in different inferential connexions, they would have different meanings in virtue of this fact.

The conception of understanding Brandom develops is an inherently *normative* one. It belongs to a Kantian tradition that distinguishes the mental as that which is the appropriate object of a certain kind of *evaluation*:

Judging and acting involve *commitments*. They are *endorsements*, exercises of *authority*. Responsibility, commitment, endorsement, authority—these are all *normative* notions. Judgements and actions make knowers and agents liable to characteristic kinds of *normative* assessment. [Brandom 2009: 32]

In making commitments, we are not *modally* bound by them; we may transgress our commitments, hold commitments that are mutually incompatible, self-defeating or the like. Sapient creatures such as ourselves have a responsibility to integrate our commitments into (what is grandly titled) a *synthetic unity of apperception*. Brandom (and Kant) uses the term ‘apperception’ in the sense found in old-fashioned works of psychology: the term means a kind of rational assimilation into a wider body of ideas, the coupling of a new strand into the web of belief. Maintaining a synthetic unity of apperception involves three responsibilities; one *critical*, one *ampliative* and one *justificatory*. The critical responsibility is to rectify mutually incompatible commitments, it is to ensure that one’s system of beliefs is *consistent*. If one maintains P , $\neg Q$ and $P \rightarrow Q$, then at least one of these commitments must be jettisoned. The ampliative responsibility is to become aware of the material consequences of one’s current commitments. Acknowledged commitments give rise to further commitments that one may not yet be aware of. The responsibility to make oneself aware of these further commitments and to integrate them appropriately into the whole is a responsibility that aims at *completeness*. Whereas the ampliative responsibility looks inferentially downstream, the justificatory responsibility looks inferentially upstream. Agents are responsible for offering reasons for their commitments, by claiming commitments that entitle them to their current commitments. The justificatory responsibility is directed at ensuring that one’s network of commitments is *warranted*.

Brandom’s explication of content and the synthetic unity of apperception provides a good way to flesh out the notions of *cohesion* and *coherence-making relationships* gestured at in the epistemological literature on understanding. It captures the features that we take to characterize understanding well. Understanding is not binary, it comes in degrees. On the inferentialist picture this is explicated in terms of the degree to which the thing understood is incorporated into one’s synthetic unity of apperception; the extent to which one acknowledges its inferential commitments, presuppositions and incompatibilities, the degree to which one has a mastery of its bearing on other propositions. As Kvanvig [2003] notes, degrees of understanding can vary along three distinct dimensions. The *breadth* of one’s understanding increases as it is embedded in a wider cognitive landscape. Broad understanding is characteristic of the philosophic enterprise. Philosophers seek to discover what the implications of results in one domain are for other domains of enquiry; what is the bearing of contemporary cosmology on the tenability of various religious commitments, of quantum mechanics on logic, or Minkowskian geometry on *a priori*? The *depth* of one’s understanding increases as it is more richly integrated into one’s web of commitments, by increasing numbers of material connexions between commitments being acknowledged or forged. Understanding can be more or less fine-grained. Finally, understanding can vary in *significance*. Two people may believe the same proposition, yet assign to it a different significance. Failure to acknowledge

the significance of a belief is to fail to grasp the way in which it is connected to other beliefs. In Bayesian contexts, different agents may both grasp a probabilistic connexion between beliefs, but assign different probabilities to this connexion. Inferentialism captures too the distinction between implicit or practical understanding, and explicit or propositional understanding. Implicit understanding of a concept involve a practical mastery of its inferential relations to other concepts. an *ability* to reliably navigate the inferential terrain. Explicit understanding involves *articulating* the material inferences that were previously merely *used* in practice. As well as making sense of these structural features of understanding, inferentialism captures the *phenomenology* of understanding. Consider the sense of untetheredness one feels when dealing with a foreign currency. The sensed need to know what you have in “real money” comes from not understanding the currency one is dealing with. If one has 600 Dirhams, this information can be embedded in one’s arithmetical inferential mastery; one knows that if they spend 100 Dirhams they will have 500 left. But what is lacking is an ability to make the kind of inferences that are normally available, given a more familiar currency: what proportion of one’s money it takes to buy a pint of milk, to catch a taxi, or to survive on for a week.

Our inferentialist conception of understanding can also be put to use to resolve some debates in the epistemological literature. The first concerns whether understanding is *factive*. Here we are required to distinguish between understanding-*of* and understanding-*why*. Understanding-*of* takes as its target *subject matters*. Algebra, quantum field theory, chess, neo-Fregeanism and the Hagiographa are all things that can be understood in this sense. Understanding these subjects involves being able to *reason* about them, in the sense adumbrated above. Understanding-*why* takes as its target *specific facts*. We might endeavour to understand why no equation of the form $x^n + y^n = z^n$ where $n > 2$ has any solutions in the positive integers, photons behave thus and so, particular configurations of chess pieces result in reciprocal zugzwang, abstraction principles could ground our arithmetical concepts, or what Qoheleth’s metaphor of vapour is intended to convey. Understanding why a proposition is true involves assimilating it into a synthetic unity of apperception, in particular by looking inferentially upstream from the proposition in order to discharge the justificatory responsibility it carries with it. Understanding-*why* is not merely a case of finding justification though—we may be justified in believing *that* something is the case, without understanding why it is the case. Even justification requires a degree of understanding; being justified in believing that something is the case involves a grasp of the kind of things that would count as evidence for the fact. Justifying *It has been raining* by alluding to the wetness of the ground, requires grasping that *It has been raining* preserves entitlement to *The ground is wet*. Proffering this justification, despite involving understanding-*of*, does not provide understanding-*why* however, since it appeals only to what is inferentially *downstream* from the proposition *It has been raining*. Understanding why it has been raining requires looking inferentially *upstream*; *viz.* knowing a proposition which preserves entitlement or commitment to *It has been raining*. Understanding-*of* is the most fundamental kind of understanding. Understanding-*why* presupposes understanding-*of* and appends additional requirements to it. Hence, understanding-*why* is a subclass of understanding.

Understanding is non-factive. Kvanvig [2003] discusses some non-factive uses of understanding, such as instances in which the term is used to hedge commitments. One may say “I understand that Greece may leave the Eurozone”. In instances such as these one uses ‘understand’ to express uncertainty; the speaker is not here committed to Greece leaving the Eurozone, but to something weaker, perhaps that they have been informed that this is so. Kvanvig rightly sets aside such uses of ‘understanding’ as not central to the sense of understanding that epistemologists should be primarily interested in. But understanding-of is clearly both central to an epistemology of understanding and non-factive in many cases. One can understand Newtonian mechanics, intuitionistic logic, phlogiston theory or astrology even though all are false. Moreover, one can understand these theories without being committed to their truth. So understanding is both non-factive and non-doxastic.⁶⁴ It may be thought that understanding-of requires some knowledge, even if this is not knowledge that the subject matter is itself accurate. A historian of science who understands phlogiston theory may not take the theory to be true, but she *does* take it to be the case that scientists really did put forward such a theory. However, whilst this may be the paradigmatic case of a person who understands phlogiston theory, even this kind of knowledge is not required for the agent to possess understanding. A historian of science may be *uninterested* in the truth or falsity of the theories she studies, and have no inkling as to whether phlogiston theory is accurate. Nevertheless, she still grasps phlogiston theory, and achieves understanding (if not *complete* understanding) of a subject matter.

The second debate in the epistemological literature concerns whether understanding requires *knowledge*. A common picture of understanding is that it is simply *a species of knowledge*, in particular knowledge of *causes*.⁶⁵ This equation of understanding with knowledge of causes is untenable however, since there are mathematical, logical, philosophical, linguistic and geometrical counterexamples to it, such as Lipton’s geometric example stated above. Lipton himself, as is clear from the example, realizes that there are explanations that are not causal explanations. His causal account of *explanation* at least is, as such, self-consciously incomplete. It may be that there are interesting things to say about causal understanding in particular, but we cannot take understanding to amount to nothing more than knowledge of causes, and so cannot on those grounds take understanding to merely be a species of knowledge. Pritchard [forthcoming] argues that understanding is not a species of knowledge because knowledge is incompatible with *environmental epistemic luck* while understanding is not. Suppose for example that a researcher at CERN makes use of an instrument to check whether she can observe superluminal neutrinos. As chance would have it, she opts to use the one instrument that is not malfunctioning. Her observations, along with some background theory, allow her to gain understanding the behaviour of

⁶⁴Contrast Kvanvig who states:

If one understands a body of information, this would seem to require knowledge of that information; and if one understands that the 2000 presidential election was at stake in Florida, it is hard to see how such understanding could obtain without knowing that the election was at state in Florida. [Kvanvig 2003: 191-2]

Indeed so. But understanding of *actual* events is itself a subclass of understanding *simpliciter*, as our examples illustrate.

⁶⁵Pritchard [forthcoming] attributes this view to Peter Achinstein [1983], Wesley Salmon [1989], Philip Kitcher [2002], James Woodward [2003], and Peter Lipton [2004]; and Grimm [forthcoming] attributes the view to Aristotle, Lewis [1986], Miller [1987], Strevens [2008] and Greco [2010, manuscript].

neutrinos, but we would not count this belief as knowledge since it is only luckily true; the researcher might easily have used a malfunctioning piece of equipment. I share Pritchard's intuitions here, although cases such as these may lie within an intuitional penumbra for many, in which their practical grasp of 'knowledge' and 'understanding' seems unable to decide the case either way. Given our inferentialist account of understanding however, it is *clear* that understanding cannot be a species of knowledge; understanding does not require truth or belief, let alone being untroubled by environmental epistemic luck. Understanding is not only not a species of knowledge, it is also more *fundamental* than knowledge. This is because understanding is required for *belief*. One cannot believe a proposition unless one grasps what the proposition *means*, at least to some extent. And this grasping is itself not possible without some practical mastery of how to *use* the proposition, *viz.* a practical mastery of the kinds of material inferential relationships it is caught up in. Understanding then is necessary for belief, and belief is not sufficient for knowledge, so understanding is more fundamental than knowledge.

Finally, our inferentialist account of understanding appears to provide the materials to resolve Meno's paradox of inquiry. Plato's Socrates states the paradox in this way:

Do you realize what a controversy you're conjuring up? The claim is that it's impossible for a man to search either for what he knows or for what he doesn't know: he wouldn't be searching for what he knows, since he knows it and that makes the search unnecessary, and he can't search for what he doesn't know either, since he doesn't even know what it is he's going to search for. [Meno, 80e; tr. Waterfield 2005: 113]

Socrates was inquiring into the nature of virtue, but any project of analysis—the analysis of knowledge, for instance—is subject to the paradox. If we already understand what knowledge is, then the analysis of knowledge appears superfluous; but, if we do not understand what knowledge is then we will have no grip on what analysis is the correct one. The inferentialist account of understanding, by providing an account of the difference between *implicit* and *explicit* understanding, offers a route through the dilemma. We come to the analytic project with an *implicit* understanding of knowledge; we know in a practical way how to *use* the term; how to apply it and how to use it to make inferences. By making explicit the inferences that a concept licenses, what is gained from the project of analysis is an *explicit* understanding of the concept of knowledge; so analysis is not superfluous. On the other hand, our implicit understanding of knowledge allows us to check it against various analyses: are there occasions where it is appropriate to ascribe knowledge but where it is not appropriate to ascribe the analysans, or vice versa? So, there is no puzzle surrounding how we can know which analysis is the correct one. Our implicit inferential mastery of concepts makes the project of analysis possible.

5.4.3 Understanding and Explanation

With an account of understanding in hand, we can now provide some minimal conditions for explanatory power. The *minimal account* of explanation is as follows:

[Min] An explanatory act is one which is aimed at producing explicit understanding in its audience.

Note the restriction to *explicit* understanding. There are of course ways of inculcating into persons a practical mastery of various tasks and even concepts without providing any kind of explicit understanding, but doing so is not explaining *per se*, apart perhaps from in some extended analogical sense. This account meets our three criteria of adequacy. The *diversity* of explanation is accounted for: explanations that are not causal, or that are not arguments are not ruled out from the outset; yet what is common to *prima facie* very different sorts of explanation, for instance conceptual analysis on the one hand and individual causal explanations on the other, is here clear. Trivially, the epistemological character of explanation, the relationship between explanation and understanding, is captured by the minimal account and not illegitimately ruled out, as it was by the D-N, S-R and unificationist models. The illocutionary nature of explanation is also made explicit here. In addition to meeting the three criteria, [Min] captures characteristic traits of explanation. Explanations can come in degrees of explanatoriness; there are more or less deep or shallow explanations. One can explain the fire by way of reference to the arsonist and the Zippo, but a fuller explanation could include facts about the combustibility of local materials, the presence of sufficient oxygen and so forth. The minimal account captures this, for the same reasons that the inferentialist account of *understanding* was able to capture the fact that understanding comes in degrees. Similarly, we can make sense of the fact that there can be different, non-overlapping but noncompeting explanations of the same event. One answer to the question *Why is there steam venting from the kettle?* could involve Gay-Lussac's law; another, one's intentions to make a cup of tea. There are multiple ways to understand a single event; and these can be *complementary* rather than in competition, so long as one is not materially incompatible with the other. As with understanding, so too with explanation: explanations are only in competition with one another when one is materially incompatible with the other. This is in contrast with Kitcher's account of explanation, in which one looks for the *single* most unifying scheme; the explanatory store that balances the largest scope with the fewest argument patterns. There, the single most unifying scheme *trumps* the others. Finally, [Min] explains the non-*transitivity* of explanation: X can explain Y, and Y can explain Z, yet X does not explain Z. This can take place because there can be material inferential relations between X and Y, and Y and Z, but not X and Z. For similar reasons, [Min] illustrates why explanatory regresses are not troubling. One can explain X in terms of Y, but have no explanation of Y itself.⁶⁶ This happens when there are relations of commitment preservation and entitlement from Y to X, but no (or few) relations of commitment preservation and entitlement from some other subject matter Z, to Y. [Min] provides the correct result—that despite the regress there is an explanation of X—because X is integrated in the right kind of way into the subject's synthetic unity of apperception in virtue of the explanatory act.

The minimal account also has the virtue of elucidating what it is that varies amongst different kinds of explanation. The difference in explaining a topic or explicating a concept, on the one hand, and explain-

⁶⁶Perhaps explaining events by appealing to laws of nature is like this.

ing *why* an event e occurred, on the other hand, is reflected in the difference between understanding *simpliciter* (i.e. understanding-of), and its narrower sibling understanding-why. Similarly, as we noted earlier, [Min] allows one to elucidate the distinction between explaining *that* e occurred, and explaining *why* e occurred. The former involves understanding of the most general sort. In order to grasp what constitutes a reason to think that e occurred, one must have some inferential grasp of e ; an ability to discern the sorts of things that, if e occurred, could or would be the case. Explaining *why* e occurred involves an ability to discern the sorts of things that could or would *make it the case that* e .⁶⁷ A correct explanation of e is the factive version of this; believing what has in fact made it the case that e .

5.4.4 The Enhanced Indispensability Argument Revisited

Let us consider again Baker's reconstruction of the explanation of the prime periods exhibited by periodic cicadas:

- (1) Having a life-cycle period which minimizes intersection with other (nearby / lower) periods is evolutionarily advantageous. [**biological 'law'**]
- (2) Prime periods minimize intersection (compared to non-prime periods). [**number theoretic theorem**]
- (3) Hence organisms with periodic life-cycles are likely to evolve periods that are prime. [**'mixed' biological / mathematical law**] [Baker 2005, p.233]

On the minimal account of explanation, this is clearly explanatory. The prime periods exhibited by cicadas in their nymphal stage were cognitive *danglers*. While biologists knew that they took place, they could not inferentially integrate them into a wider body of understanding in any interesting way. A grasp of (1) and (2) capacitates this kind of cognitive integration. The mere explanatoriness of an utterance, on our account, however *does not* license or otherwise commit one to the existence of the entities it quantifies over. All that is required is that on hearing the utterance, and grasping it, understanding is produced in the listener. An explanation can be something that *enables* understanding, without itself being true. This is perhaps clearest in the case of *metaphorical* or *analogical* explanations. Uttering "Love is a journey" enables understanding in the listener, whilst being strictly false. It is also the case in examples of scientific explanations that make use of falsehoods, such as those described by Batterman.⁶⁸ As such, platonists cannot merely appeal to the explanatoriness of mathematics in the sciences to warrant belief in mathematical objects, for uttering (false) mathematical sentences may enable understanding, just as uttering (false) metaphorical sentences does. Perhaps mathematical explanations form a special subclass of explanations, a subclass that, as a result of its own peculiar characteristics, *does* warrant belief in the entities that these explanations quantify over.

⁶⁷One must construe 'making it the case that' in the broadest terms, *viz.*, not merely causal terms. That there are two horizontal dimensions but only a single vertical one, for instance, *makes it the case* that when sticks are thrown into the air, there are, at any given moment, more near the horizontal axis than the vertical.

⁶⁸See also Elgin [2007, 2009] for examples of explanatory falsehoods in science.

This, however, is not the case. Mathematical explanations of concrete phenomena are a species of analogical explanation. One can see this by first noting that, Baker's derivation of (3) from (1) and (2) is *invalid*: (3) does not follow from (1) and (2) alone. To make the inference valid, *bridge principles* are required. The purpose of these bridge principles is to capture relevant similarities that we take to hold between the integers and the years. The number-theoretic theorem that prime periods minimize intersections explains the prime-periods, given the salient structural similarities between integers and time segmented into year-periods, we can infer (within relevant parameters) that what holds for the integers holds for years. This is a form of analogical explanation. Certain structural features of a system are explained by pointing out that it is relevantly similar to another system whose structural features are manifest to us. The prime-period life-cycles of cicadas is demystified because we come to realise that time, partitioned into years, has structural features that, in collaboration with other facts about the evolutionary pressures on cicadas and their physiology, make more likely the fact that periodic cicadas have prime-period life-cycles. Analogical explanations take place *via* representations. Recall van Fraassen's representation scheme:

Z uses X to represent Y as F

In Baker's example, biologists use inscriptions, speech acts and so on, to represent years as integers. In doing so, their inferential mastery of number theory is co-opted to allow them to reason about time, which, in combination with biological data, produces understanding. If one does not represent years as integers, one's understanding of number-theory is not so much as *relevant* to the case of periodic cicadas. The possibility of analogical explanation undermines the thought that there is an exclusive dividing line between representation and explanation. This is relevant because Baker insists that the Enhanced Indispensability Argument shows that this exclusive distinction exists, and uses it to refute positions which, as Baker sees it, do not respect the representational/explanatory division. For instance, Leng [2005] points out that it is common for scientists to knowingly make use of models that involve idealization and, hence, falsehood, but which still manage to 'represent truths about physical systems' [ibid.: 11], or, as we have parsed it, to capture the nomic or modal structure of concrete systems. The suggestion here is that, given that this practice is acceptable, one might wonder why false-because-mathematical assumptions should be thought of any differently. Baker responds:

[Leng's] argument misses the main point of the Enhanced Indispensability Argument, which is precisely to draw a sharp line between representational and explanatory uses of mathematics. [...] Leng accepts that the mathematics in the cicada case is genuinely explanatory, so her argument based on idealized concrete posits is off target. [Baker 2009, p626]

But it is Baker's objection that is wide of the mark. The mathematical explanation of the life-cycles of periodic cicadas is explanatory *because* one can represent time, partitioned into yearly chunks, as the integers. Analogical and mathematical explanations are explanatory precisely *because* they are

representational; there is no sharp line to be drawn.⁶⁹

5.5 Renewing Ontology

In attempting to unearth the presuppositions underlying the indispensability argument, I criticized the Quinian method of ontology. Whilst there is nothing *wrong* as such in criticizing a view and leaving it at that, so long as the criticism is cogent, I believe that philosophers have a constructive as well as destructive duty. There is little point in critiques that do not aim to advance our understanding. In this spirit then I will conclude the chapter with a suggestion for what should replace Quine's vision of ontology. In order to do so, I want to tie together two strands of our discussion; that which pertained to the relationship between the Quinian method of ontology and realism, and our more recent discussion of the ontological preconditions for the application of mathematics, including explanatory applications. I have criticized the Quinian method of ontology, and the critique of nominalism it proffers, on two grounds. Finding truthmakers for our best theories is not *sufficient* for realism of an interesting sort—the species of realism that involves our ability to come to know what the world is like—because it entails ontological anti-realism. Believing that the world consists only of an infinite number of wheels of liverot, is to get things wrong about the way the world is. Nor is the goal of finding truthmakers *necessary* for this kind of realism. The manner in which we gain knowledge of the world cannot in all cases be construed as endorsing an *argument*; inferring a conclusion from a set of accepted premisses. A constellation of practices can inculcate in persons an ability to *reason* correctly about the world, and hence to grasp its nomic or modal structure. Mathematical language is a vehicle of this process, but it can carry out this role without itself accurately describing how things stand in an independently existing abstract domain of reality. In this way, our theories enable us to track the structure of the world without being true. Seen in this light, the ability to come to grasp what the world is like by *making use of* false sentences should not seem mysterious. It is simply the ability—found in humans, teachable in chimpanzees, and absent in monkeys—to think *analogically*; an ability studied naturalistically by cognitive scientists.

The second ground on which I criticized the Quinian method of ontology was that in looking to what our best theory of the world quantifies over, the true locus of ontological commitment had been missed. On the present view answers to the metaphysical questions of this or that variety turn on providing an adequate understanding of our practices involving this or that area of discourse. Semantics is one part of that. The philosopher's task is to make sense of scientific and mathematical practice as it is actually conducted; to unearth and understand the conditions that make possible the application of mathematics in the sciences. The guiding question for the ontologist is not then 'What is required to exist to make our best theories true?', but 'What is required to exist in order to make sense of *what takes place*?'. This

⁶⁹Baker's criticism is a little surprising given that this is a point which Leng *herself* makes [Leng 2005: 16], and, as we noted earlier, since his own pragmatic account of explanation does not entail any such segregation of representational and explanatory uses of mathematics.

includes facts about the success of the application of mathematics and its demystifying, explanatory, role. This is the turn from a purely semantic outlook, to a more encompassing one, involving facts, not just about the semantics and syntax of scientific theories, but about their wider role in our cognitive ecology and practice. The elision of pragmatic aspects of science has acted to skew ontology in favour of platonism; for, by attending only to what is quantified over, the Quinian method of ontology tells us we are committed to whatever we need to talk about in order to carry out our best science. When we examine the *role* that mathematical discourse plays in scientific practice, it seems clear that we can usefully talk about mathematical objects without committing ourselves to their existence; that mathematical talk would serve its purpose equally well regardless of the existence of the objects it describes.

How a theory is *used* is essential to determining our ontological commitments. For whilst a sentence or theory may be said to have ontological commitments *itself*, in some slightly extended sense, what is of interest to the ontologist is not the ontological commitments of a theory itself, but rather the ontological commitments one accrues by using a theory. The question is how one is to determine *these*. In the preceding part of this chapter—by proffering an account of what is required to exist for our mathematical practices, as we find them, to take place—this is precisely what I have been attempting to do. Our goal now is to state things in more general terms; to articulate a *recipe* for ontological commitment that will replace Quine's.

5.5.1 Pragmatic Meta-Analysis

Ontological commitment, so I say, is not a straightforward function of the meaning of our best theories, but is borne out of a more complex relationship between their meaning and use. Huw Price has made a similar claim with regards to *naturalism*. Normative vocabulary may not be reducible to naturalistic vocabulary, but 'subject naturalism'—the thesis that human persons are natural creatures describable by science—requires no such reduction, only an explanation of how this vocabulary could play the role it does amongst natural creatures.⁷⁰

The challenge is ... simply to explain in naturalistic terms how creatures like us come to talk in these various ways. This is a matter of explaining what role the different language games play in our lives—what differences there are between the functions of talk of value and the functions of talk of electrons, for example. [Price 2004: 87]

Whether or not Price's project succeeds is not a topic that we will take up here, but the methodological pragmatism that he advocates and that I have urged in the case of nominalism, is of interest. We can bring into sharper focus this method of ontology by adopting the notion of a *pragmatic metavocabulary*, introduced and explicated by Robert Brandom. Exploring semantic relations between vocabularies, is part of the diurnal practice of analytic philosophers. Field's programme of nominalization falls under this rubric: the goal is to show that nominalistic vocabularies have the same expressive power, when it


⁷⁰The phrase 'natural creatures' is, of course, etymologically infelicitous for the naturalist.

comes to describing the concrete world, as mathematical vocabularies. There are however, other kinds of relations between vocabularies. Two kinds of *pragmatically mediated* semantic relations are of interest to us here, *PV-sufficiency* and *VP-sufficiency*:

PV-Sufficiency: A cluster of practices or abilities P , is PV-sufficient for a vocabulary V , when engaging in those practices or abilities is sufficient for someone to count as *deploying* the vocabulary V .

VP-Sufficiency: A vocabulary V is VP-sufficient for a cluster of practices-or-abilities P when the vocabulary is sufficient to *specify* those practices or abilities.

As such, if a vocabulary V is VP-sufficient to specify PV-sufficient practices for a vocabulary V' , then V allows one to say what one must do in order to deploy the vocabulary V' . Brandom illustrates these pragmatically mediated semantic relations in a precise way by considering a purely syntactic case. Let an *alphabet* be a finite set of primitive sign types. The *universe* of that alphabet is made up of all the finite strings of signs that can be generated by concatenating primitive signs from the alphabet. A *vocabulary* is any subset of the universe generated by the alphabet. *Deploying* a vocabulary in this syntactic sense involves an ability to ‘read’ and ‘write’ the vocabulary; *viz.* an ability to distinguish strings that belong to a specified vocabulary and an ability to produce all and only the strings that belong to a specified vocabulary. Given these definitions one can describe the abilities that are PV-sufficient to use a vocabulary in terms of *finite-state automata* (FSAs). Brandom considers the alphabet $\{a, h, o, !\}$ with the ‘laughing Santa vocabulary’ consisting of strings such as ‘hahaha!’, ‘hohoho!’, ‘hohahahoha!’ and so forth. The laughing Santa FSA can be represented by the following diagram:



The nodes 1-4 represent the *states* of the FSA and arrows *state transitions*. The laughing Santa FSA can both read and write the laughing Santa vocabulary. As a reader it must recognize only legitimate strings of the vocabulary, and it does so beginning at the first character in a string and progressing until it either stops or reaches the end. If the first character of the string is not an ‘h’, the string is rejected. If it is an ‘h’ the FSA transitions into state 2. Only if the character there is an ‘a’ or an ‘o’ will the FSA move transition into state 3, and so on. The string is recognized as legitimate just in case the FSA arrives at state 4. A similar story can be told of the laughing Santa FSA as a *writer* of the laughing Santa vocabulary. The FSA begins in state 1, produces an ‘h’ and transitions into state 2. In this state it selects at random an ‘a’ or ‘o’, produces it and transitions into state 3, and so on. When the FSA reaches state 4 it halts.

We can think of the diagraph as being a kind of vocabulary, V_δ .⁷¹ This vocabulary is *VP-sufficient* for the laughing Santa FSA; it is able to specify the ‘abilities’ possessed by the laughing Santa FSA. This cluster of abilities is, in turn, sufficient to produce the laughing Santa vocabulary; it is *PV-sufficient* for that vocabulary. Adopting Brandom’s terminology, V_δ is a *pragmatic metavocabulary* of the laughing Santa vocabulary. Moreover, it is a pragmatic metavocabulary for a vocabulary which it itself cannot employ, and so is a case of pragmatic expressive bootstrapping: the diagraph tells us everything that something needs to do to deploy the laughing Santa vocabulary, but is unable to deploy that vocabulary itself. This syntactic sense of a vocabulary provides a clean example of pragmatic metavocabularies and pragmatic expressive bootstrapping, but the more philosophically interesting cases involve *meaningful* vocabularies. Nominalistic and mathematical vocabularies are one such case. In the context in which Brandom is writing, ontology is not of primary concern. Brandom’s goal is to analyse the relationships between logical, modal, normative and intentional vocabularies. Traditional projects of analysis sought to do this by bringing to light *semantic* relations between those vocabularies. Brandom broadens the project of analysis by introducing *pragmatically mediated* semantic relations, that by doing so other kinds of relationships between vocabularies may be disclosed. As such, the practices-or-abilities Brandom is interested in are those that are required or sufficient for using certain *vocabularies*, but I want to generalise the notion here. While mathematical practices certainly involve employing mathematical vocabularies, they go beyond this. The practices we are concerned with here involve deploying mathematical vocabularies in order to describe the world, make inferences and predictions about it, and explain its features, and our goal as ontologists is to unearth what is required to exist in order for these things to take place.

Over the course of this chapter, I have been developing the case that one can provide a nominalistically acceptable explanation of the role of mathematics: one can say, in nominalistic terms, what is required to exist and to take place in order that mathematics can be successfully applied in scientific, and other, contexts. In other words, nominalistic discourse provides a *pragmatic metavocabulary* for our mathematical practices. Moreover, it is a case of *pragmatic expressive bootstrapping*: nominalistic vocabularies are expressively weaker than mathematical ones (indeed *strictly* so, since, as we noted earlier, mathematical vocabularies are not in general conservative over nominalistic ones). My claim is this: *ontological commitment is determined by the content of the pragmatic metavocabulary of our best theory of the world*. The correct project for nominalists is not to show that our best theory of the world does not quantify over mathematical objects, but that the pragmatic metavocabulary for that theory does not quantify over mathematical objects.

Quine took it that our ontological commitments could be discerned by looking at what is quantified over by our best theory of the world. This, as we saw, was problematic. The correct locus of ontological commitment is not in our best theories *themselves*, but in our best accounts *of* those theories. Quine’s

⁷¹If it seems implausible that a diagraph is a vocabulary of *any* kind, then note that the laughing Santa FSA could be represented by means of a state-table, or simply a list of instructions.

observation about the ability to permute the ontological commitments of our theories at will—from ordered n -tuples of real numbers to sets, and so on—is an illustration of this (although this is not what Quine took it to be an illustration of). The ontology of our best model of the world does not matter *because it is a model*. All that matters is that representing the model permits surrogative reasoning about its target structure. Because we represent concrete systems as mathematical structures, looking to what is quantified over by our theories to uncover our ontological commitments conflates the model with that which is being modelled. So the semantic content of our best theories is at best only an *indirect* guide to what the world is like. One can think of a pragmatic metavocabulary, on the other hand, as giving an account of what is ‘going on’ when we employ a certain vocabulary or engage in other practices and activities. If our best account of what is going on (in general) does not commit us to abstract objects then we are not committed to the existence of abstract objects. Our best descriptions and explanations of the world may be couched in a mathematical vocabulary, but when we become reflective about what we are *doing* when we employ that vocabulary it may turn out that we are not committed to mathematical objects. On reflection, this is inevitable: *merely employing* a certain vocabulary does not tell us anything about what we are committed to ontologically, we can only *begin* to answer that question when we ask what we are doing *by* employing that vocabulary; when we turn to pragmatics. (Another way of putting this: a vocabulary on its own, floating free as it were of any particular constellation of practices cannot be committed to anything; to say it could would be to make a category error. So in assessing ontological commitment, we *cannot* simply appeal to the contents of a particular vocabulary, we *must* look to the practices within which the vocabulary in question is embedded and ask *of those practices* whether we can make sense of their success and what we must suppose exists in order to make sense of their success.)

On this view, ontology is still constrained by our best science, but not in the straightforward way that the Quinian supposed. One must still look to good science as a guide to what exists. On the other hand, the practice of ontology cannot, as was promised by the Quinian method, simply be handed over to science. Teasing out the ontological commitments of our best theories by composing a pragmatic metavocabulary for them is a peculiarly *philosophical* task, although one which, so it seems to me, can and must be conducted in a proper spirit of deference to the scientific enterprise. The deployment of mathematical vocabularies within our mathematical practices—which is to say, quantification over mathematical objects in our best scientific theories of the world—does not entail commitment to mathematical objects, because a merely nominalistic vocabulary is VP-sufficient to specify those practices; it is a case of strict pragmatic expressive bootstrapping. If a nominalistic pragmatic metavocabulary can be provided for any given area of abstract discourse, then that discourse is available to the nominalist to make use of with impunity. The future of nominalism lies in discharging this goal, and in the next chapter I will sketch out the lineaments of this programme.

Chapter 6

Rolling Out the Project

In the last chapter I developed an outworking of the thought that what is quantified over, even in sentences we take to be correct, is not, apart from in an indirect way, a guide to ontological commitment. There, as in most of the literature on abstract objects, the locus of our attention was on mathematical objects. Mathematical objects however are far from the only abstract objects quantified over in technical and ordinary discourse. The issues raised by mathematical objects can be seen as a single instance of a more general problem for nominalism, which we might call *the problem of pervasiveness*. This is that talk of abstracta permeates many areas of our discourse far beyond the mathematical, and that quantification over abstract objects takes place in a good deal of sentences which we take, at least *prima facie*, to be true. For the nominalist who insists on paraphrasing our discourse to avoid quantification over abstract objects, the problem of pervasiveness is acute, as nominalization will be required not only in physics but across the many other areas of discourse that make use of quantification over myriad kinds of abstract objects. It is not obvious, or even plausible, that a paraphrasing strategy that is successful in one domain will still be so when transposed into another domain. I have argued that the shift from “reconstructive” nominalism to that which aims to make nominalistic sense of mathematicized theories without attempting to paraphrase them is not merely a *tactical* advance for nominalism, but a *philosophical* one; one that makes better and deeper sense of what the issues really are. It is also an approach that holds out the promise of providing a *unified* account of nominalism; an account that explains why *all* useful quantification over abstracta is harmless. In the proceeding I will argue that the conceptual machinery earlier employed to show that the application of mathematics in the sciences does not involve ontological commitment to mathematical objects, can also be applied to other areas of discourse in which abstract objects are quantified over; in particular *proposition-talk* and *type-talk*. While this is hardly comprehensive, it offers some credence to the claim that a unified nominalistic account of quantification over abstract objects is possible, and hence to nominalism *tout court*, as opposed to merely mathematical nominalism.

6.1 Propositions as Tools of Measurement

The ‘received view’ treats propositional attitudes as, fundamentally, relations between the bearer of the propositional attitude—a person, or her mind—and an object of the attitude—a proposition, or some other semantically evaluable object. It is certainly true that this reflects the logical form of propositional attitude ascriptions. For instance:

Ai Weiwei believes that the ruling party’s censorship laws are draconian.

Abdullah bin Abdulaziz al Saud hopes that Wahhabism will remain popular.

Bashar al-Assad fears that he will be deposed.

These all take the form $t\Psi A$, where t is a term, Ψ an intentional operator, and A a formula; *viz* they specify a relation between an agent and a proposition. Not all intentional attitudes are propositional attitudes. Intentional verbs can be predicates that take noun-phrases as their complement, for instance:

Solomon worships Yahweh.

Emeline fears the Cookie Monster.

These are both grammatically acceptable, as are some other constructions, for instance:

Alfred Pennyworth knows who Batman is.

Thom Yorke knows how to play guitar.

The compliments to the intentional verbs in these constructions are neither noun phrases nor propositions. Interestingly, few intentional verbs can sensibly take complements of this sort; the verb ‘know’ is unusual in this respect.¹ Our focus however will be on propositional attitudes. It is not just that propositions are quantified over; they also are involved in inferences that seem intuitively correct:

Žižek believes that *The Sound of Music* is implicitly anti-semitic.

Nobody else believes that *The Sound of Music* is implicitly anti-semitic.

∴ There is something that only Žižek believes.²

From the platonist’s point of view, this inferential felicitousness is a sign that one is getting things right about the way the world is when one quantifies over propositions. One can think of an intentional operator Ψ as a function f that takes as its argument persons and maps them to a proposition or set of propositions, so that $f_b(\text{Ai Weiwei}) = \langle \text{the ruling party’s censorship laws are draconian} \rangle$, where f_b is the BELIEF function, and $f_h(\text{King Abdulla}) = \langle \text{Wahhabism will remain popular} \rangle$, where f_h is the HOPE function.³ Propositional attitude attributions involve relating an agent to an abstract object that represents the attitude. This bears a striking resemblance to the manner in which measurements represent

¹Cf. Priest [2005: ch1] for an overview of these issues.

²Note the close resemblance to the kind of inferences discussed in chapter 4.2.2.

³I adopt Horwich’s [1998] convention of using $\langle p \rangle$ to denote *the proposition that p*.

objects as having certain magnitudes; *viz.* by relating the object measured to an abstract object—usually a real number or a vector—that represents the object. So, for instance, the mass-in-kg function, f_{kg} , assigns objects a numerical magnitude depending on an intrinsic physical quantity; in this case, their mass. For example, $f_{kg}(\text{Earth}) = 5.97219 \times 10^{24}$ and $f_{kg}(\text{muon}) = 1.9 \times 10^{-28}$. This has led a number of philosophers to speculate that propositional attitudes are not best understood as a relation between a person and a proposition, but rather as a feature of the agent herself, ‘measured’ by a proposition. The leading idea then of a measurement-theoretic account of propositional attitudes ascriptions is that the predicates that one uses to attribute propositional attitudes function as measure predicates. Although these predicates are relational in form, their *role* is to attribute *intrinsic* psychological states or properties to agents by locating that agent in a logical space; just as other measure predicates, although relational in form, attribute intrinsic magnitudes to physical objects by locating them in a logical space. Propositions are abstract objects which one can represent in order to *model* propositional attitudes. Robert Matthews summarizes the view thus:

Although both sorts of predicate appear to be used to ascribe to their logical subjects a relation to an abstract entity, a proposition or a number, it has been argued that appearances are deceptive: the predicates are relational in form, yet in neither case do they express genuine relations. To say that an object has a temperature of 20 degrees Celsius is not to say that the object stands in a relation to the number 20; likewise, to say that a subject has a certain propositional attitude is not to say that the subject stands in a relation to a proposition. Rather, it is to attribute to that person a certain psychological state which is specified by means of its location in a measurement space, in just the way that we specify the temperature of an object by means of its location on a measurement scale. [Matthews 1994: 131]⁴

Whilst Matthews [1994 and, esp., 2007], has done the most to develop a view of this sort, thinking of propositional attitude ascriptions in measurement-theoretic terms was first suggested by Suppes and Zinnes [1967], and expressions of the view can be found scattered throughout the philosophical literature. Paul Churchland, interestingly, sees the approach as a means of avoiding an unwanted commitment to the existence of propositions:

The idea that believing that *p* is a matter of standing in some appropriate relation to an abstract entity (the proposition that *p*) seems to me to have nothing more to recommend it than would the parallel suggestion that weighing 5 kg is at bottom a matter of standing in some suitable relation to an abstract entity (the number 5). For contexts of this latter kind, at least, the relational construal is highly procrustean. [Churchland 1979: 105]

Field also toys with the idea:

⁴Whilst this is very close to the view I believe the nominalist ought to adopt, I would add a caveat to this particular expression of it. Matthew’s claims that to ‘say that an object has a temperature of 20 degrees Celsius is not to say that the object stands in a relation to the number 20’ [my emphasis]. This may be the case, however it is not a *requirement* for the nominalist. In line with the view we have been developing, the nominalist can maintain that temperature ascriptions *do* express relations between objects and numbers, but that (i) these statements should be treated as nominalistically adequate rather than true *tout court*, and (ii) despite the logical form of these statements, what one *does* when one makes temperature ascriptions is model the object being measured—in the sense elaborated in the previous chapter—and does not thereby commit oneself to mathematical objects. Another way of putting this: appearances *are* deceptive, but *metaphysically* (as opposed to *semantically*) so.

[W]hat is explanatorily basic is not relations like *x has mass-in-kg r* between physical objects and numbers; rather, what is explanatorily basic is certain ‘intrinsic relations’ holding among physical objects themselves, e.g., *the mass of x is the sum of the masses of y and z*. ... Consequently, assigning real numbers to physical objects can be viewed as just a convenient way of discussing the intrinsic mass-relations that those objects have. [...] The obvious analog to the postulate of a rich array of intrinsic relations among physical objects in the number case is to postulate a rich array of intrinsic relations among internal states inside the believer in the psychological case [Field 2001: 70]

As may be expected, Field holds that, homologous with his programme for mathematical nominalism, propositional talk would have to be ultimately *dispensable*:

[T]he relations among these states must be ‘intrinsic’ *in the sense that they are to be describable independent of propositions*. The intrinsic relations and the assumptions about them must be powerful enough to prove a representation theorem that shows that whenever the assumptions are satisfied there is a mapping of internal states into propositions that preserves the kind of structure that is important to propositions. [ibid., my emphasis]

Davidson also has endorsed the measurement analogy. Unlike Field, who also endorses the notion that mental representation involves token *sentential* mental representations, Davidson sees measurement-theoretic accounts as a means of unpacking a non-relational and non-representational account of the propositional attitudes:

Just as in measuring weight we need a collection of entities which have a structure in which we can reflect the relations between weighty objects, so in attributing states of belief (and other propositional attitudes) we need a collection of entities related in ways that will allow us to keep track of the relevant properties of the various psychological states. [...] [W]e do not need to suppose that there are such entities as beliefs. Nor do we have to invent objects to serve as the ‘objects of beliefs’ or what is before the mind, or in the brain. For the entities we mention to help specify a state of mind do not have to play any *psychological* or epistemological role at all, just as numbers play no physical role. [Davidson 1989: 11]

Finally, in a recent book by Soames, a measurement-theoretic account of propositional attitude ascriptions is *assumed*:

[T]he role structured propositions play in our theories [is] as a measure of cognitive states in which agents predicate properties of things. In physical theory we use numbers, and other abstract objects, to talk about theoretically significant relations that physical magnitudes of various sorts bear to one another. In semantic and psychological theory we use abstract propositional structures to talk about theoretically significant relations that representational cognitive states and activities bear to one another, and to the world. The conditions the theory specifies as necessary and sufficient for entertaining these propositional structures are what allow us to use them to track the relationships that hold among actual and possible predications by agents. [Soames 2010a: 91]

Here I will sketch the measurement-theoretic account, argue that, as with the case of mathematical measurement, the practices of propositional attitude ascription make sense without presupposing the

existence of propositions, and, finally, provide some arguments that the measurement-theoretic account is indeed the correct account of propositional attitude ascriptions.⁵

6.1.1 Developing the Measurement-Theoretic Account

The state of enquiry into the nature of the mental and how (or if) mental properties are realized in physical substrates, is not such that we are in a position to fully describe propositional attitudes in terms of their neurophysiological or computational substrates. It may even be, as, for example, Daniel Dennett holds, that at the subpersonal level there are no discrete computational processes that would correspond to representations at the personal level; i.e. that propositional attitudes are not natural kinds at computational or neurological levels of description.⁶ Because of this, Matthews takes it that our conceptual grip on the propositional attitudes is *functional*, in terms of their role in behaviour. Whilst the link between the possession of propositional attitudes and behaviour is not *analytic*, as logical behaviourists took it to be, propositional attitudes do bear prominent *contingent* relations to behaviour. Because of this, attitude ascriptions can be used to predict the behaviour of agents. In the broadest possible terms, an agent who desires a certain state of affairs will act to bring it about, an agent who fears a certain state of affairs will act to avoid it, and so on. How a propositional attitude is manifested in behaviour however will depend on the numerous other propositional attitudes held by the agent, as well as other factors; perhaps, for instance, pertaining to the agent's competency in bringing about or avoiding states of affairs, or her ability to deduce what else must obtain if a given state of affairs obtained. Despite these mediating factors, propositional attitude attributions provide a generally reliable, if not infallible, means of predicting and explaining behaviour, because we are relatively skilled in presuming the kinds of background, ambient propositional attitudes that agents normally possess, and in identifying and accounting for these other mediating factors.

As Matthews notes, in these respects the manner in which one predicates propositional attitudes is not qualitatively different from a predicative scheme which is pervasive in our descriptive parlance. The scheme involves specifying the capacities, abilities and dispositions of persons and things by relating them to the entities or events that are the *object* or *target* of those capacities, abilities and dispositions:

Thus, people can be flag wavers, firefighters, stock traders, psychologists, proctologists, devil worshippers, furniture buyers, risk takers, English speakers, opera lovers, bus drivers, skirt chasers, and so on, while animals can be anteaters, nest builders, bloodsuckers, disease carriers, and so on, and inanimate objects can be can-openers, posthole diggers, earth movers, pile drivers, dishwashers, typewriters, word processors, and so on. [Matthews 2007: 151]

⁵See also Dennett [1987] and Stalnaker [1984] for endorsements of the measurement analogy. The quotations from Churchland and Davidson can be found in Matthews [2007].

⁶Another possibility, that fewer contemporary philosophers are willing to countenance, is that propositional attitudes are not supervenient upon neurophysiological substrates but are the product of some other substance, which may, or may not, be susceptible to empirical research.

Matthews calls these ‘two-dimensional’ characterizations, as they specify the capacities, abilities or dispositions of an entity along *two axes*: the first involving the object of these capacities, abilities or dispositions, and the second the relation that the entity bears to that object. Instead of relating persons to things, as is the case in basic two-dimensional characterizations, propositional attitude ascriptions related persons to propositions. Propositions have truth conditions, which is to say that specific features of the world are required to obtain in order for a proposition to be true. As a result, relating persons to propositions rather than objects permits more finely-grained characterizations with concomitantly more finely-grained predictive and explanatory powers than basic two-dimensional characterizations are capable of providing. Calling someone a firefighter may be informative as to her capacities and abilities, but has little predictive power, since some firefighters—such as those stationed on military bases—may never actually fight fires. Knowing of a particular firefighter that she believes a particular action—dousing flames with water perhaps, or sealing off a particular doorway by a particular means—would put the fire out, permits predictions with quite specific contents. If one knows that a firefighter believes that sealing off the hallway on the third floor of the building will slow the fire, and that the firefighter intends or desires to slow the fire, then one can predict that the firefighter will (attempt to) seal off the hallway on the third floor of the building. In the case of propositional attitude ascriptions then, the first axis of the two-dimensional characterization corresponds to the state of affairs that is the object of the propositional attitude, and the second axis picks out the relation that the person bears to that state of affairs. Roughly (and defeasibly) speaking, if *t* fears *A*, *t* will make efforts to ensure that *A* is not made true, if *t* desires *A*, *t* will make efforts to ensure that *A* is made true, and if *t* believes *A* then this fact will enter into *t*’s behaviour in various familiar ways. For instance, if one knows that the firefighter endorses the material inference from <the building has a good alarm system> to <the building is safe to work in> and intends that buildings be safe to work in, one can, taking ambient propositional attitudes and the like into account, predict that the firefighter will attempt to make it the case that buildings possess good alarm systems.⁷ Thus, propositional attitude ascriptions exploit both the vertical, semantic properties attributed to propositions—their conditions of satisfaction—and their inferential, horizontal relations to each other, in order to make fine-grained predictions about behaviour of other agents.

Not only do we use propositional attitudes in this way; our grip on what propositional attitudes *are* is derived from our practices of ascribing propositional attitudes to people. (It is not as though the notion of a propositional attitude was arrived at by cognitive scientists studying the neurological or computational

⁷Treating propositional attitudes in causal-functional terms is not the only option however. One can think of propositional attitude ascriptions as measuring *commitments*; ‘deontic scorekeeping’ in Brandom’s sense. Understood this way, propositional attitude ascriptions are a means of keeping track of the commitments and entitlements of other persons. Saying “Ailsa believes that the colour card is antique white” for instance, allows one to keep track of the commitments and entitlements of Ailsa by surrogatively reasoning with <the colour card is antique white>. Since <the colour card is antique white> entails <<the colour card is red>is false> and <something is antique white>, for instance, one can, by reasoning in the logical space of propositions, come to infer that Ailsa is committed to the colour card’s not being red and there being something that is antique white. Propositional attitude ascriptions becomes means of predicting behaviour *indirectly*, when they are coupled with the assumption that the person in question *acknowledges* (or fails to acknowledge) her commitments. Whilst I favour this option, for clarity of exposition I will continue to talk of propositional attitude ascriptions in causal-functional terms, rather than, at every juncture where it is pertinent, describe both options. Nothing hangs on which option is taken, as far as a defence of nominalism is concerned.

structure of the brain.) It is because of this that Matthews holds that any characterization of propositional attitudes will be functional.⁸ Physicalists will hold that these functional properties supervene on neurophysiological or computational properties, but these properties play no role in the practices of predicating propositional attitudes, and forming predictions and explanations on this basis. What is required for these practices is only that propositional attitudes possess the right sorts of relationships to the states of affairs described by the propositions used to characterize them. The *inferential* relations among propositions—or, better, a relevant subset of those inferential relations, where what is relevant depends on pragmatic factors—must track the *causal* and *constitutive* relations that hold between the propositional attitudes themselves, if reasoning about propositions is to allow surrogative reasoning about propositional attitudes. It is plausible, for example, conjunction elimination tracks a constitutive relation of many propositional attitudes: <Andy will win Wimbledon and the US Open> entails <Any will win Wimbledon>; and Judy's hope that Andy will win Wimbledon is plausibly constitutive of Judy's hope that Any will win Wimbledon and the US Open, in the sense that the effects characteristic of the former propositional attitude are a subset of the effects characteristic of the latter propositional attitude. On this basis, Matthews even provides an informal sketch of how a representation theorem might be proven.⁹ Since the propositional attitudes are characterized functionally, a detailed functional characterization of the propositional attitudes would provide an intrinsic description of them. Given this, one could characterize two structures. Firstly, $\mathfrak{R}_\Lambda = \langle \Lambda, \lambda_i^f, \vee \rangle$, where Λ is the domain of propositions, λ_i^f are the first-order n-adic relations defined on Λ and \vee is a unary function on λ_i^f that assigns to each relation in λ_i^f an extension in Λ . Secondly, $\mathfrak{R}_P = \langle I, P_i^f, \rho_i^s, \vee \rangle$, where I is the domain of persons who possess propositional attitudes, P_i^f are the first-order properties of I that are attitudes, ρ_i^s are the second-order functional properties and relations of the propositional attitudes, and \vee is a function that assigns an extension in I to every first-order property in P_i^f , and an extension in P_i^f to every second-order property and relation in ρ_i^s . Given specifications of \mathfrak{R}_Λ and \mathfrak{R}_P , a representation theorem would show that there is a mapping from the propositional attitudes P_i^f to the propositions λ_i^f that preserves the structure of the second-order relations ρ_i^s on the propositional attitudes. Because propositional attitudes are defined functionally, none of this would require a knowledge of the underlying neurophysiological or computational structures that ultimately manifest these functional properties. In addition, and relatedly, it cannot be doubted that propositional attitudes have the functional properties that propositional attitude ascriptions attribute to them, since propositional attitudes are specified *in terms* of those propositional attitude ascriptions. This is not prescriptive as to the intrinsic character of those mental states. Matthews rejects the language of thought hypothesis, and uses his measurement-theoretic account to undercut motivation for the view. This much is legitimate; if the structure of propositional attitudes themselves cannot be simply read off the logical structure of propositional attitude attributions, then an important motivation

⁸Although, as I have noted, a *deontic* characterization of propositional attitudes is also available.

⁹A representation theorem shows that a particular distinguished class of models of a theory—distinguished for being of particular conceptual interest—exemplifies, within isomorphism, *every* model of the theory. More formally, if \mathfrak{M} is the set of models of a theory and \mathfrak{B} is some conceptually distinguished subset of \mathfrak{M} , a representation theorem for \mathfrak{M} with respect to \mathfrak{B} claims that for any model M in \mathfrak{M} there is a model B in \mathfrak{B} , such that B is isomorphic to M . Representation theorems can be in terms of homomorphisms as well as isomorphisms. Cf. Suppes [2002] for discussion.

for the view that propositional attitudes *are* a relation between a person and an ‘object’ of the attitude is lost. The measurement-theoretic account is not however *incompatible* with relational views of this sort.

6.1.2 Propositions and the Pragmatics of Representation

I have provided a sketch of the measurement-theoretic account of propositional attitude ascriptions. Notably, Matthews often treats the account as though it avoids ontological commitment to propositions, although he does not elaborate on why this should be so. As it stands, the measurement-theoretic account quantifies over abstract objects, propositions; but to determine the ontological commitments of these measurement practices, we require an account of what makes possible the practice of using proposition-talk in this way—what I have been calling a pragmatic meta-theory. The applicability of proposition-talk parallels closely the applicability of mathematics. In both cases, abstract objects are used to *model* concrete systems by *measuring* them, and hence representing them as thus and so. Earlier, I characterized this in terms of representing-*as*. To say that propositional attitudes are represented *as* propositions is perhaps to stretch the terminology somewhat, but the process is analogous to the mathematical case. A person t is represented by an abstract structure $t^* = \langle b, f, d \rangle$. Here, b corresponds to the belief of t , f to the fear of t , and d to the desire of t .¹⁰ Nothing so precise as an equation governs our understanding of the interactions of different propositional attitudes in an agent, but there are rough and ready ways in which we take them to interact. Perhaps it is impossible for a person to both fear and desire the same thing; in which case this would be represented by the model by endorsing $\forall p[(f=\langle p \rangle) \rightarrow \neg(d=\langle p \rangle)]$. Principles such as *if t fears A , t will make efforts to ensure that A is not made true* and *if t desires A , t will make efforts to ensure that A is made true* correspond in the mathematical case to bridge principles linking concrete with mathematical facts. As such the same considerations that applied to the mathematical case carry over to the case of propositions. In the first place, a pragmatic meta-theory of representation offers support for the measurement-theoretic account. Earlier, I defended Suárez’s inferential conception of representation:

[Inf] X represents Y only if (i) the representational force of X points towards Y , and (ii) X allows competent and informed agents to draw specific inferences regarding Y . [Suárez 2004: 773]

Given this account of representation, it is clear that propositions are used to represent propositional attitudes. (i) is met, as the aim of relating agents to propositions is to ascribe properties to these agents, and (ii) is met, as doing so allows competent and informed agents to draw inferences regarding the behaviours of the people they attribute propositional attitudes to. Using a logical space of propositions to reason surrogatively about mental states *just is* to use them as a measure of mental states. Matthews

¹⁰This is of course a significant idealization, normal agents have many beliefs, fears and desires, and many other propositional attitudes besides belief, fear and desire. A less idealized model would be $t^* = \langle \{a_1^1, \dots, a_n^1\}, \{a_1^2, \dots, a_n^2\}, \dots, \{a_1^n, \dots, a_n^n\} \rangle$, where each a^i represents a different kind of propositional attitude.

characterizes the use of propositions to model propositional attitudes in terms of a structure-preserving mapping from the propositional attitudes to the propositions. As with the mathematical case, this appears, on the face of it, to involve a commitment to abstract objects, since for a mapping between the propositions and the propositional attitudes to *exist*, there must *be* propositions. However, as with the mathematical case, to construe representation as a dyadic, metaphysical relation between the source and the target of representation, is an idealization, and one which misrepresents the ontological baggage of our representational practices. Recall van Fraassen's meliorated representation schema:

Z uses X to represent Y as F

Propositions are not literally used to represent propositional attitudes. Rather, persons use thoughts, inscriptions and utterances to represent propositional attitudes *as* propositions. Whether propositions exist is, as with the mathematical case, orthogonal to this practice. Matthews' own formal account of how propositions represent propositional attitudes should play the same role in our thinking as Bueno, Colyvan and French's formal accounts of mathematical representation. These, I argued, did not provide an account of how mathematics can be used in a representational way, as *no* metaphysical relations of the sort they considered were necessary or sufficient for representation to take place. These formal accounts did however contain an important vestige of truth: they make explicit the structural features that mathematical objects would have if they existed, and in doing so explain why reasoning about those mathematical objects allows for surrogative reasoning about the concrete systems they are used to represent. Formal accounts of the representation of propositional attitudes by means of propositions, such as that sketched by Matthews, also explain why reasoning about propositions allows for surrogative reasoning about propositional attitudes, by making explicit the structural features that propositions would have if they existed.

As with the mathematical case, there is a naturalistic story to be told about the way in which this kind of representation provides epistemic access to the world. Understanding phenolphthalein consists in a mastery of the inferences that one ought to make about phenolphthalein, an ability to reason about the substance. Similarly, understanding propositional attitudes consists in an ability to reason about propositional attitudes. By representing propositional attitudes as propositions, reasoning about propositions can facilitate reasoning about propositional attitudes, and thus facilitate understanding of propositional attitudes. Our ability to reason analogically is then sufficient to explain why reasoning about propositions allows us to reason about propositional attitudes, and hence how the practice provides epistemic access to the world. What is required for the successful application of proposition-talk is not *knowledge* of the features of an abstract domain of propositions, but an *understanding* of what the features of such a domain would be, were it to exist. Finally, nominalists ought to account for the intuitive correctness of inferences such as:

Žižek believes that *The Sound of Music* is implicitly anti-semitic.

Nobody else believes that *The Sound of Music* is implicitly anti-semitic.

∴ There is something that only Žižek believes.

in terms of their being *nominalistically adequate*. While they are strictly false, the premisses are nominalistically adequate in that the concrete world is the way it would have to be for them to be true. Making inferences of this sort take one from nominalistically adequate premisses to nominalistically adequate conclusions. As such, in most contexts—*viz.*, non-philosophical contexts—asserting these claims and endorsing these inferences is pragmatically felicitous.

6.1.3 Must Propositional Attitudes be Relations?

We have noted that propositional attitude ascriptions are of the logical form $t\Psi A$, and, as such, express relations between agents and propositions. As Matthews notes, this has been pivotal in shaping the philosophical landscape:

Proponents of the Received View are convinced of the truth of the relational conception, even in the absence of empirical and abductive support for the Received View, because they presume that the relational character of the attitudes can simply be read off the sentences by which we canonically attribute propositional attitudes. [Matthews 2007: 98]

Nominalism is characterized by a scepticism about our ability to straightforwardly read the structure of reality off the logical form of the sentences one uses to describe it; as such, merely *noting* the relation form of propositional attitude ascriptions is not (or *ought* not to be) particularly convincing line for the relationist to take. In this section however I want to consider firstly some *arguments* from the logical form of attitude ascriptions that aim to demonstrate that intentional attitudes *must* be relational in nature. Following this, I will make the case that propositional attitude ascriptions must be measurement-theoretic in nature.

Priest and the Relational View

Graham Priest has raised some interesting arguments for the view that intentional attitudes must be relational that we will now address. While these apply specifically to intentional *predicates* rather than operators, which are the focus of our attention, the two are sufficiently affiliated that, were there good reasons to suppose intentional attitudes towards objects were inherently relational, these considerations would in all likelihood carry over to intentional attitudes towards propositions. As such, the arguments merit separate consideration. Consider the attitude ascription:

Benny fears the man next door.

Instead of taking this to involve, fundamentally, a relation between Benny and the man next door, one might understand the ascription as positing an intrinsic property of Benny, perhaps, as Priest suggests, that ‘Benny has a man-next-door-representation in his mental “fear box”’ [Priest 2005: 58]. Priest presents what he takes to be a problem for this understanding of intentional attitudes:

[I]t makes nonsense of ‘There is someone whom Benny fears, but who is, in fact, a very nice man.’ This is of the form: $\exists x(bF'x \wedge Mx)$. To make sense of the first conjunct, the quantifier has to range over representations, but then the second conjunct is nonsense: the representation is not a man at all—nice or otherwise. [ibid.]¹¹

One can avoid this consequence by introducing a representation relation Rxy into the language, meaning ‘ x is a representation of y ’, and parsing ‘There is someone whom Benny fears, but who is, in fact, a very nice man’ as $\exists x\exists y(bF'x \wedge Rxy \wedge My)$. This however is itself problematic. In the first place:

[H]ow is one to understand ‘Benny and Penny fear something (the same thing)’? $\exists x(bF'x \wedge pF'x)$ won’t do: there is no guarantee that Benny and Penny have exactly the same mental representation of the object in question. [ibid.]

One could define an equivalence relation \sim between representations, such that $x \sim y$ if and only if x and y are representations that appear to be of the same thing. The problem, as Priest sees it, is in how to define \sim , as different representations of the same thing can be arbitrarily different. In the second place, this still involves quantifying over the object of the intentional attitude, which, as Priest points out, may not exist.

The measurement-theoretic account of propositional attitudes however is well-placed to handle such objections. Notice that Priest understands the claim that intentional attitudes do not fundamentally or metaphysically involve a relation between an agent and the object of her belief, as *semantic* or *hermeneutic* in character. If intentional attitudes do not fundamentally involve a relation between an agent and an external object then this feature of the structure of reality must be reflected in the logical structure of our intentional attitude ascriptions; hence the introduction of the representation relation Rxy . But this is precisely what the measurement-theoretic account disclaims. As such, in the formalization $\exists x(bF'x \wedge Mx)$ of ‘There is someone whom Benny fears, but who is, in fact, a very nice man’ it need not be presumed that the quantifier must range over representations in the first conjunct. The quantifier ranges over people in both conjuncts; this is simply an instance of the two-dimensional predicative scheme that permeates our descriptive language.¹² Ascriptions such as these may of course quantify over non-existent objects, but the nominalist, in parallel to the case of measurement involving real numbers, vectors and other mathematical objects, can treat these ascriptions as nominalistically adequate rather than true *tout court*. Nor do attributions of the same intentional attitude to different people cause any special problems for the measurement-theoretic account. As propositional and other intentional attitudes are defined functionally, there is no requirement that two individuals must represent an object in the same way if they are to be correctly ascribed the same intentional attitude. To state ‘Benny and Penny fear something (the same thing)’ is to attribute to both a certain propensity or disposition. Even this need not (and in the ordinary case *will not*) require endorsing the same counterfactuals with respect to Benny and Penny’s behaviour.

¹¹Recall that Priest uses $\exists x$ to denote the existential quantifier.

¹²Notice however that this *is* a problem for those that hold that the structure of intentional attitudes *can* be read off the logical structure of intentional attitude ascriptions. Either they must take the relation to be between an agent and a mental representation, in which case Priest’s objections apply, or they must take the relation to be between an agent and the object of her thought, in which case they must countenance non-existent objects, as Priest does.

This will depend on what other propositional attitudes one ascribes to Benny and Penny, as well as other mediating factors involving Benny and Penny's distinct abilities, competencies and the like.

6.1.4 Are Measurement-Theoretic Accounts Inevitable?

One serious problem with the view that propositional attitudes involve standing in particular relations to propositions involves the very abstractness of propositions:

[I]t is difficult to see how the mere standing in a relation to an abstract entity ... could possibly be causally efficacious, which beliefs surely are. Believing that there are still two beers in the fridge may prompt me to ask you to stay for another beer, but what gets my lips and tongue moving to issue this invitation is not some relation that I bear to an abstract entity, specifically a proposition, but some physical goings on in my head that are capable of causing these physical events. [Matthews 2007: 105]

Since propositions are not causally efficacious, merely standing in a relation to a proposition cannot itself account for the causal efficacy of propositional attitudes. This parallels our earlier discussion of heavy-duty platonism. There it was noted that, for example, a system with a temperature of 40°C has different counterfactual properties to a similar system with a temperature of 30°C. To adopt Balaguer's idiom, some purely concrete fact must be holding up its end of the propositional attitude bargain. One option here is to retain the relational picture of propositional attitudes, but to abandon the notion that what one is related to when one has a propositional attitude is an *abstract* object. For instance, the relation could be a quasi-linguistic mental representation in a language of thought. This is nominalistically acceptable, but faces problems of its own. For example, we treat inferences of the following sort as valid:

Qoheleth believed that there is nothing new under the sun.

It is true that there is nothing new under the sun.

∴ Qoheleth believed something true.

However, what is denoted by the *that*-clause in the second sentence is not a mental representation, since its correctness is not dependent on the existence of mental representations. As such, if the *that*-clause contained in the first sentence denotes a mental representation, then the argument is invalid on account of involving an equivocation. Here, there appears to be something of a dilemma. On the one hand, if propositional attitudes are constituted by a relation between an agent and an abstract proposition, then their causal efficacy is ruled out. On the other hand, if propositional attitudes are constituted by a relation between an agent and a mental representation *and this is what makes propositional attitude ascriptions true*, then the inferences that are treated as valid in fact involve an equivocation. Some authors (Crimmins and Perry [1989], Richard [1990] and Crimmins [1992]) have responded by, as it were, grasping both horns of the dilemma, by positing *two* different objects of belief; a semantic

object, which accounts for the validity of inferences such as the one above, and a psychological object, which accounts for the causal efficacy of propositional attitudes. The details of these accounts need not concern us here; what is salient is that some explanation must be given of the *relevance* of the semantic object of propositional attitudes to the psychological states of the agents to which they apply.

[D]istinguishing between semantic and psychological objects of belief comes with a price ... [O]ne then has to explain how the former manages to track the latter, and how native speakers are able to exploit the fact that the former tracks the latter to obtain information about the believer. For otherwise one has no account of how true belief ascriptions manage to be informative. [Matthews 2007: 115]

If the semantic object of belief is distinct from the psychological object of belief, then an account of what one is doing when one uses propositional attitude ascriptions that take abstract propositions as their complement—as in the inference above—must make sense of how this practice is informative, how it is efficacious. Some account must be given of how this practice can be used to explain the behaviour of and make predictions about the actions of agents. However, if the semantic object of belief is distinct from the psychological object then to do this requires an explanation of the relationship between the two; i.e. an account of the manner in which the semantic object of belief can be said to *track* the psychological object of belief. This task however, *is precisely what the measurement-theoretic account of propositional attitude ascriptions sets out to do*.¹³ Perhaps it is the case that all roads lead to measurement theory.

6.2 Types and Tokens

A *token* is a particular, concrete instance of a *type*, and a type is a general thing of which there can be particular, concrete tokens. Linda Wetzel [2006] uses the elegant example of a line from the Gertrude Stein poem *Sacred Emily* to illustrate the distinction:

Rose is a rose is a rose is a rose.

How many words are in this line, ten or three? The question is ambiguous between using ‘word’ to denote a word type and to denote a word token. There are ten tokenings of words in the line, but only three different types of words: ‘rose’, ‘is’ and ‘a’. One could also appeal to the longest known grammatical ‘single word’ sentence in English:

Buffalo buffalo Buffalo buffalo buffalo buffalo Buffalo buffalo.

The sentence pertains to buffalos (bison) from Buffalo (“Buffalo buffalos”) who buffalo (bully) other buffalos from Buffalo.¹⁴ Here there are eight word tokens but only one word type. Types are abstract

¹³With only one dissimilarity: the measurement-theoretic account does not assume from the outset that propositional attitudes involve a relation between an agent and a mental representation (although the view is not precluded).

¹⁴Buffalo bison [whom] Buffalo bison bully [then] bully Buffalo bison. [Forsyth 2011]

objects; the type ‘rose’, unlike its particular written or spoken instances, has no particular spatial location and is causally inert.¹⁵ In general, the kind of properties that can coherently be imputed to types are not properties that can be coherently imputed to concrete tokens. Wetzel provides the following example:

[C]onsider the grizzly (or brown) bear, *Ursus arctos horribilis*. At one time its U.S. range was most of the area west of the Missouri River, and it numbered 10,000 in California alone. Today its U.S. range is Montana, Wyoming, and Idaho, and it numbers fewer than 1,000. Of course no particular bear numbers 1,000, and no particular bear has ever has a range comprising most of the area west of the Missouri. It is a *type* of bear, a species of bear, that has both properties. [Wetzel 2009: xi]

Wetzel [2000, 2006, 2009] has defended platonism about types, and, concomitantly, argued that type nominalism is untenable. This is an important innovation in a scholarly literature that has hitherto oriented itself predominantly around the topic of mathematical objects. It has, as we noted earlier, highlighted the pervasiveness of quantification over abstract objects in everyday and scientific discourse, and provides a new challenge to any comprehensive defence of nominalism. I however will argue that nominalism about types is not only tenable, but that it makes better sense of our type-practices than does platonism. I will provide an outline of Wetzel’s case for type platonism—which essentially involves an *indispensability* argument for types—before raising some problems for the platonist account, and, finally, defending the nominalist account, of the kind provided for mathematical objects and propositions, against the main objections Wetzel levels against it.

6.2.1 Instances of Type-Talk

Wetzel begins by making a convincing case that type-talk thoroughly permeates our language. Although type-talk crops up in a number of different areas of discourse—talk of theories, works of art and so forth—I will take two representative samples, one from linguistics the other from biology—as two disciplines where type-talk plays a particularly crucial role—for purposes of illustration. The first is from the the paper ‘Language as Available Sound: Phonetics’ from *An Encyclopedia of Language* (reference to and quantification over types has been italicized):

The nature of *the syllable* has been ... a matter for considerable discussion and debate ... [M]ost native speakers of a *language* can recognise the *syllables* of their own *language* ... Various *hypotheses* have been suggested: that *the syllable* is either a *unit* which contains an *auditorily prominent element*, or a *physiological unit* based on respiratory activity, or a *neurophysiological unit* in the *speech programming mechanism*. *The concept of the syllable* as a *phonological*, as distinct from a *phonetic*, *unit* is less controversial [Collinge 1990: 8, in Wetzel 2009: 7]

¹⁵Note that a puzzle attaches to both these examples. The *lines* themselves are word types. Consider the line from *Sacred Emily* which consists of ten words. These cannot be word tokens, because tokens are concrete and the line (type) is abstract. But nor can they be types, since there are only three word types in the line. Wetzel’s theory of word types handles this paradox by distinguishing between tokens and *occurrences*. This is an issue however which we can bracket.

Here *the syllable* and *the speech programming mechanism* are quantified over. Both are presumably types of which there can be individual, concrete instances. Similarly, from environmental biology,

The ivory-billed woodpecker, once North-America's largest and most spectacular, was declared extinct. Less than a century ago, *it* was found across the South. *Its* last confirmed sighting in the United States was in Louisiana in the 1950s.

The banded bog skimmer, a rare dragonfly, was found for the first time in Maine. *It* was the first time *it* had appeared so far north.

The Tarahumara frog, which lived in Arizona, has disappeared from the United States. However, *it* is still found in Sonora, Mexico. [Dicke 1996: B12, in Wetzel 2009: 11]

Here there is apparent reference to the types *the ivory-billed woodpecker*, *the banded bog skimmer* and *the Tarahumara frog*. Note that the properties attributed to them could not be attributed to individual, concrete tokens: individual ivory-billed woodpeckers are not declared extinct; individual banded bog skimmers cannot be rare; and individual Tarahumara frogs have not disappeared from the United States—at least, this is not what is communicated by the passage above. Having established the near ubiquity of type-talk in everyday and scientific discourse,¹⁶ Wetzel argues, adopting Quinian and Fregean criteria for ontological commitment, that, without the existence of types, sentences that refer to or quantify over them will be false. Here, as before, I will not take issue with this point, as I hold that the existential quantifier does in fact express existence. Instead, and in accordance with the rest of this thesis, I will argue that the nominalist can and should understand sentences that quantify over abstract types as literally false, but that this is not at odds with either the ability to use these sentences to convey information about the world, nor with the usefulness of type-talk for various purposes.

6.2.2 Difficulties with Paraphrase Strategies

Although Wetzel's monograph *Types & Tokens* was published in 2009, it was largely written before non-reconstructive—or, less pejoratively, *non-paraphrastic*—defences of nominalism had become prominent in the philosophical literature. As such, Wetzel assumes that any tenable nominalism will treat type-talk as a *façon de parler* for a claim which is *not about* types, and her criticisms of nominalism largely center around the difficulty of producing a strategy for paraphrasing type-talk into sentences that only quantify over concreta. Whilst I will defend a non-paraphrastic form of type nominalism, it will be instructive for discussions later in the chapter to survey the various attempts at paraphrase strategies. Wetzel considers six different paraphrase strategies for sentences of the form 'The *T* is *P*', where *t* is a token of type *T*, and finds them all wanting:

1. Every *t* is *P*.
2. Every normal *t* is *P*.

¹⁶If readers are unconvinced of this point by the sample selection above, they are advised to evaluate Wetzel's own more comprehensive survey.

3. Most *ts* are *P*.
4. Average *ts* are *P*.
5. The disjunction of 1–4.
6. *ts* are *P* (where this is a generic or 'characterizing statement').

1 is the strategy of Goodman [1977] and Sellars [1963]. There are however counterexamples. In the biological case, for instance, it is correct to say that 'the grizzly bear is brown' but it is false that all grizzly bears are brown, some are black or blond. One might take it upon herself to shave a grizzly bear, thus rendering it non-brown, but it would remain true that the grizzly bear is brown. It fails too in the linguistic case. The word 'colour' is spelled both $c^{\wedge}o^{\wedge}l^{\wedge}o^{\wedge}u^{\wedge}r$ and $c^{\wedge}o^{\wedge}l^{\wedge}o^{\wedge}r$, but it is not true (in fact impossible) that every token of 'colour' is spelled both $c^{\wedge}o^{\wedge}l^{\wedge}o^{\wedge}u^{\wedge}r$ and $c^{\wedge}o^{\wedge}l^{\wedge}o^{\wedge}r$. Words are particularly problematic for this paraphrase strategy, as a single word can be written or spoken, have multiple spellings, senses and pronunciations. Strategies 1-5 are in fact undercut in one fell swoop. For consider the example of Wetzel's; a type-sentence she takes to be true:

The loggerhead turtle lives at least thirty years and may live fifty years.

But it is neither the case that every loggerhead turtle lives at least thirty years, that every normal loggerhead lives at least thirty years, that most loggerheads live at least thirty years, or that average loggerheads live at least thirty years. Most loggerhead turtles die before they reach the age of two.

What of strategy 6? Can one paraphrase all sentences of the form 'The *T* is *P*' into sentences of the form '*ts* are *P*'? Wetzel calls these 'characterizing statements', but in line with the practices of linguists, I will call them 'generics'. Generics come in three syntactic forms:

- *Bare Plural (BS) Generics*: e.g. Bears are ferocious.
- *Indefinite Singular (IS) Generics*: e.g. A bear is ferocious.
- *Definite Singular (DS) Generics*: e.g. The bear is ferocious

So in fact, linguists treat type-sentences as a kind of generic; a definite singular (DS) generic. As such, one might hold out hope that DS generics could, in all cases, be replaced with BS generics. Wetzel however rejects this possibility on the grounds that the kind of properties that can coherently be predicated of types cannot be coherently predicated of individual tokens. Consider again the grizzly bear, *Ursus horribilis*, whose range is Montana, Wyoming, and Idaho, and numbers fewer than 1000. These are not attributes that can be imputed to individual objects, nor can, for instance, *being extinct* or *being rare*. Perhaps Wetzel is a little too quick here, as these properties seemingly *can* be attributed using BS generics, for example:

Dinosaurs are extinct.

Snow leopards are rare.

Grizzly bears range over Montana, Wyoming and Idaho.

Grizzly bears number fewer than 1000.

DS generics *do* behave differently to BS generics however, as the former are often less well suited to express generalizations that can be expressed in BS generic form, for instance:

Shops sell goods.

The shop sells goods.*

While the latter is perfectly well constructed, it does not seem to express the *generalization* inherent in the BS generic that precedes it; rather we take it to pertain to a *particular* shop.¹⁷ What the type platonist requires however for the objection to go through is that there are instances where BS generics are *less* well suited than DS generics at expressing properties that apply to collections or pluralities of things. Strategy 6 is *prima facie* plausible, if it is to be undermined then, it is incumbent upon the type platonist to show that there are indeed instances of this sort.

6.2.3 Problems for Type Platonism

While my goal here is primarily to defend type nominalism, I will try to cause some problems for type platonism along the way; not only because this will highlight the comparative benefits of nominalism, but also because doing so will provide important hints as to what types *are*—or would be, if they existed—and hence to the role of type-talk in our cognitive ecology.

Is Type Platonism Arbitrary?

The first problem with type platonism is that it appears to introduce a degree of ontological *arbitrariness*. Recall the rejection of Goodman and Sellars' strategy for paraphrasing away type-talk. It was shown not to be the case that 'the *T* is *P*' is true if and only if every *t* is *P*. This entails that there is nothing that all and only tokens of a particular type have in common, no property that they all share.¹⁸ As a result there is often a degree of arbitrariness in where the boundaries are drawn between different types. This can be illustrated with the notion of a *species*. Biologists often individuate species *morphologically*, *viz.* in terms of *resemblance*. The morphological approach to the individuation of species—which

¹⁷Strangely enough, in a recent survey [Leslie, Khemlani, Prasada & Glucksberg 2009] it was found that participants did not, with any statistical significance, find DS generics less natural ways of expressing generalizations than BS generics, even in cases such as that above. This point however, does not harm the nominalist's case.

¹⁸Wetzel states:

[T]here is nothing interesting all and only uttered tokens of a particular word have in common other than being tokens of the word, and ... tokens of a word make up a real kind, just as members of a species do. Presumably, then, I am committed to the claim that there is nothing interesting (known or unknown) that all and only members of a living species have in common other than being members of that species (i.e.. no nontrivial, interesting, "natural," projectible property). [Wetzel 2009: 106-7]

Wetzel endorses¹⁹—takes a species to be a set of individuals that resemble each other. The resemblance however is not parsed in terms of strict necessary and sufficient conditions for species membership; if it were then the Goodman-Sellars strategy for paraphrasing away type-talk would be successful. Instead, species individuation relies on a looser notion of family resemblance.²⁰ Boll weevils resemble each other in being six millimeters long, and feeding on cotton buds and flowers, but not all boll weevils need have these properties. Moreover, the species *Anthonomus grandis*, to give the creature its Latin name, belongs to the genus *Anthonomus*, the family *Curculionidae*, the order *Coleoptera*, the class *Insecta*, the phylum *Arthropoda* and the kingdom *Animalia*, but the borders of these categories are themselves not fixed by necessary and sufficient conditions; a degree of arbitrariness or choice is involved. What is true of the biological case, carries over to the linguistic case:

The vocabulary of a widely-diffused and highly-cultivated living language ... may be compared to one of those natural groups of the zoologist or botanist, wherein typical species forming the characteristic nucleus of the order are linked on every side to other species, in which the typical character is less and less distinctly apparent. ... *For the convenience of classification*, the naturalist may draw the line ... but Nature has drawn it nowhere. So the English Vocabulary contains a nucleus or central mass of many thousand words whose “Anglicity” is unquestioned. ... But they are linked on every side with other words which are less and less entitled to this appellation. ... Yet *practical utility has some bounds*, and a Dictionary has definite limits: the lexicographer must, like the naturalist, “draw the line somewhere,” in each diverging direction. [Murray 1971, cited in Wetzel 2009: 111, my emphasis]

Wetzel summarizes the commitments of the type platonist in light of these facts:

[*W*]hat is important about a word token is what type it is a token of. Word types are pigeonholes [sic] by means of which we classify tokens; a word type alone unifies *all* its tokens in absence of any observable natural “similarity” that all the tokens have. [Wetzel 2009: 112]

We can generalize this consideration in the following way. Consider a series of particular objects $a_1, a_2, a_3, a_4, \dots, a_n$, such that each a_{i+1} resembles a_i , but where there is an object a_{i+k} , such that $i+k < n$, which does not resemble a_i . Each object resembles the object that falls before it in the sequence, but since resemblance relations are non-transitive, and because there are enough objects in the sequence, at least two objects in the sequence do not resemble each other. Whenever this is the case, there will be a degree of arbitrariness where the boundaries of types lie. This arbitrariness is to be expected if, as I will argue, type-talk is merely a nominalistically adequate means of categorizing parts of the world in useful ways. If type-talk merely reflects a practice of categorization, one would expect a degree of arbitrariness and choice, determined by contextual and other pragmatic considerations. Where no more conceptually

¹⁹Genetic and lineage approaches are both ultimately parasitic on the morphological approach. The former because it requires a notion of genetic similarity, and the latter because it ultimately requires a *synchronic* taxonomy, which itself will be parsed in terms of similarity. The population approach, which takes a species to be an interbreeding population that is reproductively isolated from other populations, faces counterexamples.

²⁰As Wetzel notes: “[I]nvariably there is diversity among members of the kinds, ... none of the characteristic properties is had by all members of the kind, and ... each of the properties may be had by members of other kinds. [Wetzel 2009: 107]”

or nomologically principled carvings are forthcoming, things in the world can be categorized according to family resemblances; no strict necessary and sufficient conditions for membership of a category are required for the practice to be useful.

For such arbitrariness to be a feature of *nature*—or a feature of the relations between nature and a non-natural domain of platonic types—is less palatable. The type realist must hold that *there is a fact of the matter* about which tokens bear the metaphysical relationship of *being a token of* to a particular abstract type. If a token is a token of a particular type, some account of *why* this is the case, is wanted. This is forthcoming for the nominalist; there are historical, pragmatic, stipulative and conventional reasons why communities might choose to categorize things one way rather than another, where there is a degree of choice or arbitrariness involved. The type realist must assert that Nature has chosen, as it were, to related certain concrete particulars to certain types—‘a word type alone unifies *all* its tokens’—even when this relation could have just as easily been otherwise, given the physical facts. The type platonist could avoid this arbitrariness by opting for metaphysical vagueness instead: there are fuzzy boundaries regarding which particulars are tokens of which type, but this fuzziness is *worldly* rather than semantic or epistemic. For many, this cure will be worse than the disease. This is hardly a refutation of type platonism, unless one comes to the table with some fairly hefty metaphysical background premisses, but it *is* an embarrassment; one which the nominalist will cheerfully circumvent.

Is Type Platonism Coherent?

A more damaging objection to type platonism can be made. Recall the following example:

The loggerhead turtle lives at least thirty years and may live fifty years.

As Wetzel pointed out, it is not the case that every loggerhead lives at least thirty years, that normal loggerheads live at least thirty years, that most loggerheads live at least thirty years, or indeed that average loggerheads live at least thirty years, as loggerhead turtles usually die before they reach the age of two. Given this, however, the following sentence is correct:

The loggerhead turtle usually dies before it reaches the age of two.

Type platonists should presumably accept the claim as being true. If one takes *The loggerhead turtle* to refer to a real object, and given that usually dying before the age of two entails not living at least thirty years, then one is committed to:

$$\exists x[Lx \wedge \forall y(Ly \rightarrow y = x) \wedge Tx]$$

where L stands for ‘is a loggerhead turtle type’, and T stands for ‘lives for at least thirty years’, and

$$\exists x[Lx \wedge \forall y(Ly \rightarrow y = x) \wedge \neg Tx]$$

which are clearly mutually inconsistent. There is a second, related objection. Type talk is, in some sense, *transparent*; one *looks through* types to the collections or pluralities they represent. This is what happens when one states, for instance, that ‘The grizzly bear once numbered 10,000 in California alone’ or that ‘The grizzly bear ranges over Montana, Wyoming and Idaho’. The problem arises when one ceases to use type-talk transparently, and looks at the type *itself*. Types are of course abstract objects. As such, the type realist is committed to the truth of:

The grizzly bear is causally inefficacious.

The existence of the grizzly bear is something we know of through its indispensable role in theory.

and, for that matter:

The grizzly bear is an abstract object.

All of these would be reckoned false by the ursinologist. As well as being committed to false propositions regarding *the grizzly bear*—and any other type we might think of—the type realist seems forced to reject seemingly *correct* propositions regarding types, for instance:

The grizzly bear kills on average four Americans every year.

The grizzly bear can be seen (by the naked eye) in Arizona.

and, for that matter:

The grizzly bear is a physical object.

The recipe for generating problematic sentences is this: either (i) predicate properties that only abstract objects possess to types that pertain to concrete objects, or (ii) predicate properties that only concrete objects possess using type-sentences. (i) produces sentences that must be true according to the type realist, but that would be regarded as false by the disciplines (such as biology or physics) that make use of such type-sentences, whilst (ii) produces sentences that would be regarded as true by the disciplines that make use of them, but must be false according to the type realist. The type realist may respond by suggesting that the problematic type sentences could be paraphrased away. This response however would be unpromising. The ease with which problematic type sentences can be generated and hence the sheer *amount* that would need to be accounted for, makes the paraphrase strategy a difficult option for the type realist. Worse still, as with the nominalist who insists that strictly false type talk must be paraphrased away, to make a convincing case that this can be done it is not enough that a handful of the problematic type-sentences are dealt with; some systematic account of how to paraphrase them away would be required.²¹ Since for every type we refer to, such problematic cases can be generated, the

²¹ Compare Wetzel, who in critiquing the nominalist states:

[It is not] enough to provide a *piecemeal* translation, by providing the odd paraphrase here and there. Given the ubiquity of apparent references to types ... we need to be assured that they can always be eliminated. *There is no assurance of this unless there is a systematic way to eliminate them.* But if their eliminativity is to be more than just a nominalist article of faith, we need to be given the elimination rules. [Wetzel 2009: 54]

type realist would have to produce elimination rules for type-talk. The nominalist then has a *tu quoque* retort. If the type realist is to avoid self-refutation, she must either accept that false type talk can be unproblematic even when it cannot be paraphrased away, or she must produce rules for paraphrasing away type-talk, that the nominalist could then purloin for her own uses.

To respond by attempting to paraphrase away problematic type-talk would be to mislocate the force of the objection however, which is this: type-talk is not *used* to characterize an abstract domain of types; it is not the *purpose* of type-talk to do such a thing. That is why, in ordinary practice—including scientific practice—we are happy to assert type sentences that violate the tenets of type realism. The purpose of type-talk is to make explicit our practices of classification, categorization and taxonomization; to codify the way in which we label worldly objects.

6.2.4 What is a Type?

In order to see why type-talk behaves in the peculiar way we have just seen, and to begin articulating a nominalist account of type-talk, we must first provide an answer to the question of what types are, and how they come to possess the properties they have. There are various ways in which the properties of a type can be extrapolated from the properties of concrete particulars. Krifka *et al.* [1995] provide (and Wetzel reiterates) seven means by which types are characterized in terms of their concrete instances. The first is the *collective property interpretation*. This involves pooling the properties of individual tokens; for instance when it is said that the grizzly bear has a U.S. range of Montana, Wyoming and Idaho, even though no individual bear has such a range. The second is the *average property interpretation*; for instance when it is said that the British family has 1.8 children. The third is the *characterizing property interpretation* which is normative in character and admits exceptions; for instance when it is said that the loggerhead turtle lives at least thirty years. The fourth is the *distinguishing property interpretation*; for instance when it is said that the Russian excels at chess. Most Russians do not in fact excel at chess, although Russia distinguishes itself from other nations by producing a large number of unusually good chess players. The fifth is the *representative object interpretation* in which an individual token can come to represent its entire kind; for example it might be said that we witnessed the loggerhead turtle in its natural habitat. The sixth is the *avant-garde interpretation*. If some token of a type possesses some exceptional property, then the property is imputed to the type itself; for instance when it is said that man set foot on the moon in 1969. Finally, there is the *internal comparison interpretation* which involves a comparison of tokens of a type along a particular parameter, where this comparison is represented as a change in the type itself. For instance, ‘The Chinese man has become significantly wealthier since the adoption of capitalism by the ruling party’.

This dovetails well with the ways in which other linguists have suggested that generics are characterized. Sarah-Jane Leslie [2007, 2008] for instance has argued that there are three kinds of generics. *Characteristic* generics are determined normatively, by the kind of thing something is. ‘Humans are

bipedal’ for instance, or ‘Lions have manes’. Characteristic generics can be correct, even when the trait only applies to a statistical minority of the creature, or artifact in question. *Majority* generics, as the name suggests, are determined by features that a majority of the concrete instances possess; ‘Barns are red’, for example. *Striking* generics, finally, are generics in which striking or unusual properties of the concrete particulars are predicated, even when these are statistically very rare amongst the concrete particulars; for instance ‘Ticks carry Lyme disease’.

What could account for the manner in which properties are imputed to types via their tokens, as well as the incoherences that a platonist view of types suffers from? The answer is this: types function as *data models* of collections or pluralities of concrete particulars. The manner in which we impute properties to types, as outlined by Krifka and Leslie, coincide with the way in which data models are constructed; as Carsten Held notes:

Consider what has to be done when constructing a data model from raw data. The first step is to eliminate outliers from the data set. They are considered the result of erroneous observation or otherwise untypical for the set. The second step is to extrapolate from the data to what could have been but was not observed. [Held 2007: 146]

Moreover, type talk is used to describe the concrete world, but it does so indirectly, mediated by an abstract type that is somehow representative of the concrete domain. So too are data models used to represent collections or pluralities of objects or events:

Finally and most importantly, the singular generic can be read such that there is direct reference to an abstract entity, but this entity itself represents the single cases. On the other hand, every data model worthy of its name models something, i.e. is a representation of the data. So, in a straightforward sense the abstract entity mediating reference to real entities in a singular generic can be identified with a data model. [Held 2007: 147]

Types then are data models: they are constructed in the same manner, and play the same representative role. This also accounts for the peculiarities we noted above. From the facts about particular loggerhead turtles, it appeared possible to accept both:

The loggerhead turtle lives at least thirty years and may live fifty years.

and

The loggerhead turtle usually dies before it reaches the age of two.

Since the type platonist accepts apparent reference to types using definite descriptions at face value, there is, for the type platonist, exactly one abstract type of which its instances are manifestations. As such, the platonist must either reject one of these correct type-sentences or accept an incoherence in their ontological commitments. For the type nominalist, the possibility of these seemingly incompatible statements is not in the least perplexing; for data models can be constructed in various different ways, using various different bridge principles, depending on the particular needs of the modeller. Here, we simply have two different data models that highlight different features of the plurality of loggerhead

turtles. In a similar vein, the type nominalist can accommodate the incorrect type-sentences that the type platonist is committed to, and can accept the correct type-sentences that are excluded by the platonist. Consider firstly:

The grizzly bear is an abstract object.

To assert this is to take the model too seriously, to conflate the model with that which it is a model of.²² Grizzly bears are represented *as* a data model, which is itself an abstract object. But one cannot simply read off the properties of a system from that which is a model of it; this depends on how the model is *used*, the manner in which it is embedded in a network of practices. A data model represents a collection or plurality of objects or events *insofar as it allows one to surrogatively make inferences* about that collection or plurality. The data model itself is an abstract object, but it does not permit the inference to the proposition that any particular grizzly bear is an abstract object:²³ Data models pool information about concrete particulars, but there is no information about particular grizzly bears being abstract objects. The correct sentences that must be treated as false by the platonist can be explained along the same lines:

The grizzly bear is a physical object.

The collected data on grizzly bears treats them all as physical objects. Although the data model is abstract, it never permits the inference to the proposition that grizzly bears *themselves* are abstract.

6.2.5 Type Nominalism

Types are *models* of (pluralities of) concrete objects. But in this case there is no special problem of types for the nominalist; they can be given a parallel treatment to both mathematical objects and propositions. *Persons use type-talk to represent pluralities of concrete entities as data models.* Because there are bridge principles—albeit often loose and intuitive ones—linking these data models to concreta, they can be used *representationally* to permit surrogative reasoning about concrete objects. As before, representing a system X as something else Y, does not involve ontological commitment to Y, as the fruits of this practice can be explained in terms of *understanding* Y—*viz.*, being able to *reason* about Ys—rather than in terms of *knowledge* of Y.

A “Dismissal” of Biology and Physics?

The question now is whether type nominalism, particularly the non-paraphrastic kind, is tenable. I hold that it is, but Wetzel has raised objections to a nominalism of this sort which must be addressed if the

²²This is the same error of which Quine was guilty in proclaiming the ‘brave new ontology’ of pure sets. By reading ontology off the best theory (i.e. *model*) of the world, Quine’s method of ontological commitment conflates the domain of the real world with the domain of the model. Using our best *pragmatic meta-theory* to determine ontological commitment avoids this error.

²³Quine makes a related error when he claims that the ‘brave new ontology’ is a pure ontology of sets. Sets may be sufficient to *model* the concrete world (and hence the only thing that is quantified over in a theory of the world), but it does not follow that the world *consists* only of sets.

view is to be defended. The first objection is that the nominalist who does not reconstruct type-talk in nominalistic terms is committed to rejecting as false the disciplines that make use of type locutions:

[A] philosophically interesting denial that the data [involving quantification over types] are *true* is either a claim that (i) the data are not, *strictly speaking*, true, but are just a manner of speaking, a *façon de parler*, for something that is strictly speaking true; or (ii) the theories of biology, linguistics, physics, and nearly all disciplines are fictions. [Wetzel 2009: 29]

Wetzel goes on to say:

Option (ii) so violates the premises of the present context, which attempts to do justice to the scientific claims made in biology and physics rather than dismiss them, that it will not be discussed. [ibid.]

This I think obscures more than it illumines, as saying that, for instance, biology is a “fiction”, is apt to make it appear on a par with a Lord Peter Wimsey mystery. Perhaps a clearer way of stating Wetzel’s claim would be something of the sort:

If type-talk is false, then sciences that quantify over types are not a good guide to the nature and structure of the world.

But how plausible is this claim? Again, I think that the notion of *nominalistic adequacy* allows the non-reconstructive nominalist to provide an account of how one can use claims which, although strictly false, can nevertheless accurately describe the world. As before, one can explicate this notion by making use of a two-sorted language $\mathcal{L}(\langle C_i \rangle, \langle M_i \rangle, \langle A_i \rangle)$ which ranges over:

(1) *concrete entities*, using primary variables x_1, x_2, \dots, x_n .

(2) *abstract entities* (objects in the embellishment), using secondary variables y_1, y_2, \dots, y_n .

and contains three kinds of predicate:

(i) *concrete predicates*, expressing relations between concreta: $\langle C_i \rangle = (C_1, C_2, \dots)$

(ii) *abstract predicates*, expressing relations between abstracta: $\langle A_i \rangle = (A_1, A_2, \dots)$

(iii) *mixed predicates*, expressing relations between concrete and abstract objects: $\langle M_i \rangle = (M_1, M_2, \dots)$

Claims about types are correct if and only if they get things right about the concrete domain, that is, if and only if they would be true at a world concretely like ours but embellished with types. By making claims about abstract objects—along with holding certain bridge principles, either explicitly or implicitly, that link them with concrete objects—one can place restrictions on the way the concrete world is, and thereby describe it indirectly. Relating concrete objects to abstract ones allows us to reason about concrete objects, to make inferences about them, and hence (and this is really the same thing) to represent them as being thus and so. This is what happens when one states that there are three dents in her car, that the mass of her car is 1000kg, or that the force required to accelerate the car at 2ms^{-2} is

2000 newtons. So too with types. The properties of types supervene on the properties of their tokens.²⁴ Since what is correct of types requires particular things to hold with regards to concrete particulars, treating type-sentences as (potentially) *correct* or *nominalistically adequate*, even if strictly false, does not amount to a rejection of the disciplines that make use of type-talk. On the other hand, because types are abstract, the *truth* of type-sentences cannot be subject to *empirical* investigation; for any empirical evidence for a claim about types would remain the same if those types did not exist.²⁵ What *can* be subject to empirical investigation is the *correctness* of type-statements. This means that despite the falsehood of all type-sentences, they are not subject to uniform rejection. Some can be confirmed as nominalistically adequate, others as nominalistically inadequate. The type nominalist *can* distinguish the correct ‘The Thylacine became extinct in 1936’ from the incorrect ‘The Thylacine became extinct in 1926’, just as the dent nominalist can distinguish the correct ‘There are three dents in my car’ from the incorrect ‘There are $10^{10^{10}}$ dents in my car’.

The Existence of Types and the Usefulness of Type-Talk

Nominalism then does not constitute a “dismissal” of disciplines that make use of type-talk in any interesting sense. There is however a second objection to nominalism to be found in Wetzel’s work. Types *themselves*, claims Wetzel, play a role in our use of type-talk. That is, the existence of types is essential to understand the efficacy of our type-discourse. This is essentially the claim that a pragmatic meta-theory of type-discourse will be committed to the existence of types.²⁶ How though does Wetzel motivate this claim, and can the nominalist counter it?

Consider a person who points to a token inscription of the word ‘colour’, and says “this is pronounced [kɔːləɪr]”. Wetzel sets up a dilemma for the nominalist. Either the referent of “this” is the token inscription of ‘colour’ or it is the word type ‘colour’. If the latter, then the type is ‘clearly not superfluous’.²⁷ If the former, then an account of how this could be informative is required. For all one has here done is point to a token inscription of ‘colour’ and uttered another spoken token of ‘colour’. Why should such a process be informative?

The answer must be that insofar as the printed word token can be said to have a pronunciation it is in the sense that it is a token of a type that has that pronunciation (even though the type may have it because some of its uttered tokens do). The type affords commonality for the spoken and written tokens together. So once again the type is not superfluous. Similarly, a spoken word token, which is an utterance, has a spelling only insofar as it is a token of a type that has that spelling. If an utterance has a spelling and a printed word a pronunciation,

²⁴Wetzel states: ‘Biology is an empirical science. Facts about species are ultimately rooted in facts about members of them—they supervene on facts about members, if you will. ... Linguistics, too, is an empirical science. Facts about words, or languages, have roots in facts about tokens of them.’ [Wetzel 2009: 83]

²⁵This is one consequence of the result that beliefs about abstract objects are *epistemically insensitive*.

²⁶It is interesting that Wetzel appears to see a need to motivate a claim of this sort, beyond the mere indispensability of type-talk. Her concern is that, without such a claim, our knowledge of words is always one-directional, from tokens to types, in which case ‘types might appear to be superfluous’ [Wetzel 2009: 121].

²⁷This is something that the anti-paraphrastic nominalist should take issue with, but it can be bracketed here.

then tokens borrow some of their properties from their respective types. And in that case *the inferential relation between types and tokens is bidirectional*: not only does the type have some properties because the tokens do, *but the tokens have some properties because the type does*. That is, the type mediates certain properties, getting them from some tokens, and conferring them on other tokens. [Wetzel 2009: 121, my italics]

The process could not be informative *sans* the existence of types, as it is the type *itself* that confers properties on tokens; the type itself confers the property of being spelled $c^{\wedge}o^{\wedge}l^{\wedge}o^{\wedge}u^{\wedge}r$ to the spoken word that is pronounced $[k\text{D}'l\text{ə}r]$, and, conversely, confers the property of being pronounced $[k\text{D}'l\text{ə}r]$ to the written word that is spelled $c^{\wedge}o^{\wedge}l^{\wedge}o^{\wedge}u^{\wedge}r$. A pragmatic meta-theory—an account of why the practices of type-talk are able to achieve what they do—will be committed to types.

The nominalist however—and indeed everyone else, as I see it—should reject this account, as what is required for this process is *not* the existence of types nor a metaphysical relation obtaining between the type and these utterances and inscriptions, whereby the type imbues its intrinsic properties of BEING PRO-NOUNCED $[k\text{D}'l\text{ə}r]$ and BEING SPELLED $c^{\wedge}o^{\wedge}l^{\wedge}o^{\wedge}u^{\wedge}r$ on them. Type-talk reflects our tendency to represent things or pluralities as belonging to kinds or classes; as *labelling* objects in particular ways. Type-talk allows one to *codify* her implicit practices of classification and taxonomization; to make explicit in what one *says* something that, in the absence of type and other consanguineous locutions, one could only *do*. What is required for the above communication to take place is merely that the speakers both classify the inscription of 'colour' and the utterance $[k\text{D}'l\text{ə}r]$ in the same way. This does not require the existence of types, only an *ability* on the part of the speakers; to classify objects of their experience in a particular way.²⁸

²⁸In fact, Wetzel *herself* goes on to outline just such a deflationary account of this exchange: 'In order to communicate what she wants via a token, characteristically a speaker must assume that her hearer associates a certain meaning with the type. She may then exploit that association in various ways.' [ibid.].

Chapter 7

Conclusion

7.1 An Autobiographical Note

My first encounter with the debate between nominalists and platonists was as an honours-level undergraduate taking a course that provided an overview of the central topics in metaphysics. Week one covered mathematical objects. As preparation for the tutorial I read Field's article 'Mathematical Objectivity and Mathematical Objects' from his collection *Truth and the Absence of Fact*. Two things struck me. Until this point, I had not been aware that *anyone* believed that numbers and other mathematical and abstract objects existed, and, moreover, I found the revelation that some people did quite startling. The second was that the means by which Field suggested one ought to discharge any commitment to mathematical objects—by demonstrating that one can, in principle, carry out scientific disciplines without talking about mathematical objects—seemed to me clearly wrong, to mislocate the issues somehow, even if I could articulate well why this would be so. So my nominalism is a *pretheoretical* one. While this is not something that fits in any obvious way into making a *case* for nominalism, it does, I think, have some epistemological import. Nominalists are often portrayed as being in the grip of a philosophical dogma, of failing to moderate their enthusiasm for philosophical argument with the emollient of pretheoretical common sense, of being too quick to see what is really an instance of *modus tollens* as an instance of *modus ponens*. However, for the *pretheoretical* nominalist, nominalism *is* common sense, and it is platonism that advocates a revision of our ordinary picture of reality on the basis of philosophical argument. So when one assumes a degree of modesty as to the ability of armchair explication and argument to warrantably alter our view of the world, pretheoretical platonists will take it to favour their view, whilst pretheoretical nominalists will take it to favour theirs. I have never really been able to shake my pretheoretical nominalism; whilst I have at points doubted the cogency of aspects of the case for nominalism I have never been able to doubt nominalism itself, and I take it that much the same is true of the pretheoretical platonist. For her, nominalism is, if a doxastic possibility at all, one that is

rather remote. Nevertheless, *having* the debate is worthwhile; for in doing so we harrow our presuppositions, refine our thinking and deepen our understanding. Understanding is the *fruit* of disagreement and doubt, indeed it is *made possible* by them: in the pursuit of reflective transcendence, *resolution* begets sterility. In this way, philosophy progresses. Metaphysical disputes in the philosophy of mathematics are no exception. Nominalism itself has transformed from the *prescriptive* programmes of intuitionism and Hilbertian formalism, to the paraphrastic programmes of Field, Chihara and Hellman, and now to the non-paraphrastic nominalism of Leng and Azzouni. We are, I believe, currently undergoing an important change in the pursuit of ontology, as our philosophical picture of mathematical practice is increasingly enriched by pragmatic considerations. We no longer seek to make alterations to mathematical practice. What has been transformed by philosophy is not mathematics or our mathematicized theories of the concrete world, but rather our understanding of them.

7.2 A Nominalist's Credo

We began with Lewis' credo, took up his challenge to account for the going concern of mathematics, and expanded it into the domains of propositions and types. Let us turn to it once again:

Renouncing classes means rejecting mathematics. That will not do. Mathematics is an established, going concern. Philosophy is as shaky as can be. To reject mathematics on philosophical grounds would be absurd. If philosophers are sorely puzzled by the classes that constitute mathematical reality, that's our problem. We shouldn't expect mathematics to go away and make our lives easier. Even if we reject mathematics gently—explaining how it can be a most useful fiction, good without being true—we still reject it, and that's still absurd. [...] I laugh at how presumptuous it would be to reject mathematics for philosophical reasons. How would you like to go and tell the mathematicians that they must change their ways, and abjure countless errors, now that philosophy has discovered that there are no classes? Will you tell them, with a straight face, to follow philosophical argument wherever it leads? If they challenge your credentials, will you boast of philosophy's other great discoveries: that motion is impossible, that a being than which no greater can be conceived cannot be conceived not to exist, that it is unthinkable that anything exists outside the mind, ... and so on, and on ad nauseam? Not me! [Lewis 1991, reiterated in 1993]

In the introduction I suggested one reason that philosophers have been reluctant to adopt nominalism: the debate is sometimes framed in terms of an *opposition* between philosophy and mathematics, or the empirical sciences. Philosophy, it is thought, is sure to lose. Another, very much related, point that I would like to address here, is that nominalism is sometimes characterized, as it is explicitly in Lewis' credo, as a *rejection* of mathematics. But this is a position that requires *philosophical* presuppositions of its own. In particular, Lewis holds that taking mathematics to be 'good without being true' constitutes a rejection of the discipline 'for philosophical reasons', but what are the philosophical reasons that underpin this claim? Presumably that mathematics is *aimed* at truth; i.e. at providing a true description

of a particular domain of abstract objects. We might call this view *teleological platonism*. Teleological platonism is a distinctly philosophical position, as opposed to a mathematical one—what mathematical axioms could such a view be derived from?—but is there philosophical *support* for the view? Recently there has been a recognition amongst philosophers that the claim that truth is the sole goal of enquiry may be a needlessly restrictive doctrine; that a more generous view of the *telos* of cognition may be in order:

[E]pistemology is the study of certain aspects of our cognitive endeavors. In particular, it aims to investigate *successful* cognition. Within its purview, then, are various kinds of cognizing, including processes such as thinking, inquiring, and reasoning; events such as changes in ones world view or the adoption of a different perspective on things; and states such as beliefs, assumptions, presuppositions, tenets, working hypotheses, and the like. Also within its purview is the variety of cognitive successes, including true beliefs and opinions, viewpoints that make sense of the course of experience, tenets that are empirically adequate, knowledge, understanding, theoretical wisdom, rational presuppositions, justified assumptions, working hypotheses likely to be true, responsible enquiry, and the like. [Kvanvig 2005: 286]

Truth is not the only epistemic goal, nor the only measure of cognitive success. Nominalism, as we have noted, only constitutes a rejection of mathematics—or the other practices that involve quantification over abstracta—when combined with *teleological platonism*. Teleological platonism however is itself a philosophical supposition, which should be made explicit so that it can itself be placed under rational scrutiny. How much rational scrutiny though can the doctrine withstand? In a survey of the methods of pure mathematics, Leng [2010: ch4] concludes that ‘nothing in our pure mathematical practices requires that we view pure mathematicians as involved in anything more than searching for axioms that characterize interesting mathematical concepts, and inquiring into what does and does not follow from the assumption that there are objects satisfying those concepts.’ [ibid.: 99] On the one hand, results are proven from axioms. On the other hand, axioms themselves are chosen. Proof from mathematical axioms clearly is not the sort of process that could establish the existence of mathematical objects. In choosing axioms, intuitions are relied on, but such intuitions, whilst pinning down the nature of our mathematical concepts, do nothing to suggest that those concepts are *satisfied*. I would suggest that the goal of mathematics is not knowledge of how things stand with a domain of abstract objects, but rather a particularly useful kind of *understanding*. The same can be said of theories of propositions and of types. I have argued the case that, in all three domains, understanding is sufficient to account for the uses to which they are put; that *knowledge* in these domains *even if it was to be had* (and it isn’t), would be strictly irrelevant to their utility.

Lewis’ rejection of nominalism is based on the Procrustean philosophical assumption that truth is the only measure of cognitive success, or at least that it is the only measure of cognitive success in the practice of mathematics. It is an assumption that appears to find little credence from the practices of pure or, as I have argued, applied mathematics. But, importantly, it *is* a philosophical assumption.

We cannot avoid philosophical suppositions, even if we can avoid philosophy. As we are, after all, entangled in philosophical commitments whatever we do, the question is not whether we should choose between our philosophical commitments and the practices that involve talk of abstracta, but rather *which* philosophical commitments make the best sense of these practices. The success of mathematics may have required abstract objects, but it did not. Making inferences about the concrete world may have required abstract objects, but it did not. Applying mathematical methods to describe and explain the world might have required abstract objects, but it did not. The practices of proposition-talk and type-talk might have required abstract objects, but they did not either. Should we accept a philosophy of those disciplines that posits abstract objects when such is a superfluity of metaphysical speculation? When it makes utterly mysterious how these fields—which are, after all, an established, going concern—could furnish us with knowledge? Not me!

Bibliography

Achinstein, P. (1983), *The Nature of Explanation*, OUP.

Adamson P. and Taylor, R. (eds) (2005) *The Cambridge Companion to Arabic Philosophy*, Cambridge University Press.

Arntzenius, F., and Dorr, C. , (2012), 'Calculus as Geometry,' in Arntzenius, F., *Space, Time, and Stuff*, OUP: Oxford.

Azzouni, J. (1997), 'Applied Mathematics, Existential Commitment and the Quine-Putnam Indispensability Thesis', *Philosophia Mathematica* 5: 193-209.

— (2004), *Deflating Existential Consequence*. OUP.

— (2006), *Tracking Reason*. OUP

— (2007), 'Ontological Commitment in the Vernacular', *Nous* 41:2, pp.204-206.

— (2009), 'Evading Truth Commitments: The Problem Reanalyzed', *Logique & Analyse* 206: 139-176

— (2010), 'Ontology and the Word 'Exist': Uneasy Relations', *Philosophia Mathematica* 18: 74-101.

Bach, K. (2005), 'Context *ex Machina*', in Szabó, Z., (ed.) (2005).

Baird, D. (2004), 'Thing knowledge: Outline of a materialist theory of knowledge', in Radder (2003).

Baker, A. (2003), 'Does the existence of mathematical objects make a difference?', *Australasian Journal of Philosophy* 81:2, pp. 246-264.

— (2005), 'Are there Genuine Mathematical Explanations of Physical Phenomena?', *Mind*, 114: 454, pp.223-238.

— (2009), 'Mathematical Explanation in Science', *BJPS* 60: 611-633.

Balaguer, M. (1996), 'A Fictionalist Account of the Indispensable Applications of Mathematics,' *Philosophical Studies*, 83, pp. 291-314.

— (1998), *Platonism and anti-platonism in mathematics*. New York: Oxford University Press.

Bangu, S. I. (2008), 'Inference to the best explanation and mathematical realism', *Synthese* 160, 13-20.

- Batterman, R.** (2002), *The Devil in the Details: Asymptotic Reasoning in Explanation, Reduction, and Emergence*, *Oxford Studies in Philosophy of Science*, Oxford: OUP.
- (2005), 'Critical phenomena and breaking drops: Innite idealizations in physics', *Studies in History and Philosophy of Modern Physics*, 36: 225-44.
- (2010), 'On the Explanatory Role of Mathematics in Empirical Science', *BJPS*, 61: 1-25.
- Benacerraf, P.** (1973) 'Mathematical Truth', *The Journal of Philosophy*, 70: 19, Seventieth Annual Meeting of the American Philosophical Association Eastern Division, pp. 661-679.
- Benacerraf, P. & H. Putnam (Eds.)** (1984), *Philosophy of mathematics: Selected Readings*, New York: Cambridge University Press.
- Bergmann, M.** (2005), 'Defeaters and Higher-Level Requirements', *Philosophical Quarterly*, 55: 419.
- Boolos, G.** (1987), 'A Curious Inference', *Journal of Philosophical Logic*, 16, pp. 1-12.
- Bowdle, B. & Gentner, D.** (2005), 'The Career of Metaphor', *Psychological Review*, 112:1, pp.193-216.
- Brandom, R.** (1994), *Making it Explicit*, Harvard University Press.
- (2000), *Articulating Reasons: An Introduction to Inferentialism*, Harvard University Press.
- (2008), *Between Saying and Doing: Towards an Analytic Pragmatism*, OUP: Oxford.
- (2009), *Reason in Philosophy*, Harvard University Press.
- Brody, B. (1972)**, 'Towards an Aristotelean Theory of Scientific Explanation', *Philosophy of Science* 39: 20-31.
- Brown, J. R.** (1985), 'Explaining the Success of Science', In Curd, M. and Cover, J. A. (eds.) *Philosophy of Science: The Central Issues*. W.V. Norton & Company, Inc.
- Bueno, O. & Colyvan, M.** (n.d.), 'An Inferential Conception of the Application of Mathematics', *Noûs*, 45(2): 345-374.
- Bueno, O., French, S. and Ladyman, J.** (2002), 'On Representing the Relationship between the Mathematical and the Empirical', *Philosophy of Science*, 69: 497-518.
- Burgess, J.** (2001), 'Review: Platonism and Anti-Platonism in Mathematics', *The Philosophical Review*, 110: 1.
- (2004), 'Mathematics and Bleak House', in Burgess, J. (2008).
- (2008), *Mathematics, Models, and Modality: Selected Philosophical Essays*, Cambridge University Press.
- Burgess, J. & Rosen, G.** (1997), *A Subject With No Object*. Clarendon Press.

- Carnap, R.** (1950), 'Empiricism, Semantics, and Ontology', *Revue Internationale de Philosophie*, 4, pp. 20-40.
- Cartwright, N.** (1983), *How the Laws of Physics Lie*. Oxford University Press.
- Cole, J. & Shapiro, S.** Review: The Indispensability of Mathematics, *Mind*, 112: 446, pp.3316.
- Chalmers, D.** (2002) 'Does conceivability entail possibility?' In *Conceivability and Possibility*, T. Gendler & J. Hawthorne, (eds.) Oxford University Press.
- (2006), 'The foundations of two-dimensional semantics', In (M. Garcia-Carpintero & J. Macia, (eds), *Two-Dimensional Semantics: Foundations and Applications*, Oxford: Oxford University Press.
- (2009) 'Ontological Anti-Realism', in Sider, T., Hawthorne, J. and Zimmerman, D (eds.) (2008).
- (2010) 'The Nature of Epistemic Space', in *Epistemic Modality*, Egan, A. and Weatherson, B. (eds.)
- Chesterton, G.K.** (1950), *The Common Man*, Sheed & Ward.
- Cheyne, C.** (1998), 'Existence Claims and Causality', *Australasian Journal of Philosophy* 76: 1, pp. 34-47.
- (2001), *Knowledge, Cause and Abstract Objects: Causal Objections to Platonism*. Springer.
- Chihara, C.** (1990), *Constructibility and Mathematical Existence*, Clarendon Press.
- Churchland, P.** (1979) *Scientific Realism and the Plasticity of Mind*, Cambridge: Cambridge University Press.
- Collinge, N. (ed.)** (1990), *An Encyclopedia of Language*, London: Routledge.
- Colyvan, M.** (2001), *The indispensability of mathematics*. New York: Oxford University Press.
- (2002), 'Mathematics and Aesthetic Considerations in Science', *Mind*, 111: 441, pp.6974.
- (2007), 'Mathematical recreation versus mathematical knowledge,' in Leng, M., Paseau, A., and Potter, M. (eds.), *Mathematical Knowledge*, OUP: New York.
- (2010), 'There is No Easy Road to Nominalism', *Mind*, 119:474, pp.285-306.
- Colyvan, M. & Lyon, A.** (2008) 'The Explanatory Power of Phase Spaces', *Philosophia Mathematica* 16: 2, 227-43.
- Comesaña, J.** (2005), 'Unsafe Knowledge', *Synthese*, 146, pp.395-404.
- Contessa, G.** (2007), 'Scientific Representation, Interpretation, and Surrogate Reasoning', *Philosophy of Science* 74: 48-68.
- (2010), 'Scientific Models and Fictional Objects', *Synthese*, 172: 215-229.

- (2011), 'Scientific Models and Representation', in Saatsi, J. and French, G. (eds.) (2011).
- da Costa, N. C. A. and French, S.** (2003), *Science and Partial Truth*, New York: OUP.
- Crimmins, M.** (1992) *Talk about Belief*, Cambridge, MA: MIT Press.
- Crimmins, M. and Perry, J.** (1989) 'The Prince and the Phone Booth: Reporting Puzzling Beliefs', *The Journal of Philosophy* 86: 685-711.
- Daly, C. and Langford, S.** (2009), 'Mathematical Explanation and Indispensability Arguments', *The Philosophical Quarterly*, 59(237): 641-658.
- Davidson, D.** (1989) 'What is Present to the Mind', in J. Brandl and W. Gombocz (eds.), *The Mind of Donald Davidson* (Amsterdam: Editions Rodopi), 3 18.
- Denet, D.** (1987) 'Beyond Belief', in *The Intentional Stance* (Cambridge, MA: MIT Press), 116 211.
- Dicke, W.** (1996), 'Numerous U.S. Plant and Freshwater Species Found in Peril', *New York Times*, January 2, p.B12.
- Dieveney, P.** (2007), 'Dispensability in the Indispensability Argument', *Synthese*, 157, pp.105-128.
- Dorr, C.** (2005) 'What We Disagree About When We Disagree About Ontology.' In *Fictionalism in Metaphysics*, ed. M. Kalderon. Oxford: OUP.
- Dretske, F.** (1971), 'Conclusive Reasons', *Australasian Journal of Philosophy* 49, pp.1-22.
- (1988), *Explaining Behaviour*, MIT Press.
- (1995), *Naturalizing the Mind*, MIT Press.
- Ducheyne, S.** (2006). 'Lessons from Galileo: The Pragmatic Model of Shared Characteristics of Scientific Representation', *Philosophia Naturalis*, 43(1): 214-234.
- Dummett, M.** (1973), *Frege: Philosophy of Language*, Gerald Duckworth & Co.
- (2010), *The Nature and Future of Philosophy*, Columbia University Press.
- Edwards, A.** (1972), *Likelihood*, Cambridge: Cambridge University Press.
- Elgin, C.** (2007), 'Exemplification, Idealization, and Scientific Understanding' in Suárez (ed.) (2007).
- (2009), 'Is Understanding Factive?', in Haddock, A., Millar, A., Pritchard, D. (eds) *Epistemic Value*, OUP.
- Everet, A. and Hofweber, T. (eds.)** (2000), *Empty Names, Fiction and the Puzzles of Non-Existence*, CSLI Publications.
- Field, H.** (1980), *Science Without Numbers*. Princeton, New Jersey: Princeton University Press.
- (1989), *Realism, Mathematics and Modality*. Wiley Blackwell.

- (1991), 'Metalogic and modality', *Philosophical Studies*, 62:1, 1-22.
- (1993). *The Conceptual Contingency of Mathematical Objects*, *Mind*, Vol 102, 406: 285-299.
- (2001), *Truth and the Absence of Fact*, Clarendon Press.
- Fine, K.** (2000) 'The Question of Realism', *Philosophers' Imprint* 1: 1-30
- (2009) 'The Question of Ontology', in *Metametaphysics*, ed. Chalmers, D., Manley, D. and Wasserman, R.: OUP.
- Frege, G.** *Grundgesetze der Arithmetik*.
- French, S.** (2003), 'A Model-Theoretic Account of Representation (or, I Don't Know Much About Art... but I Know It Involves Isomorphism)', *Philosophy of Science* 70(5): 1472-1483.
- (2010). 'Keeping Quiet on the Ontology of Models', *Synthese* 172(2).
- French, S. and Bueno, O.** (2011), 'How Theories Represent', *BJPS* 62(4): 857-894.
- French, S. and Ladyman, J.** (1999), 'Reinating the Semantic Approach', *International Studies in the Philosophy of Science*, 13: 10321.
- Friedman, M.** , 1974, 'Explanation and Scientific Understanding', *Journal of Philosophy*, 71: 5-19.
- Frigg, R.** (2006), 'Scientific Representation and the Semantic View of Theories', *Theoria* 55: 49-65.
- (2010a), 'Fiction and Scientific Representation', in Frigg, R. and Hunter, M. (eds.) (2010).
- (2010b), 'Models and Fiction', *Synthese* 172(2): 251-268.
- (2010c), 'Fiction in Science', in Woods, J. (ed.) (2010).
- Frigg, R. and Hunter, M. (eds.)** *Beyond Mimesis and Nominalism: Representation in Art and Science*, Berlin and New York: Springer.
- Garavaso, P.** (2005), 'On Frege's alleged indispensability argument', *Philosophia Mathematica* 13, pp.160-73.
- Gendler, T. and Hawthorne, J. (eds.)** (2002). *Conceivability and Possibility*, Oxford University Press.
- Gentner, D** (1983), 'Structure Mapping: A theoretical framework for analogy', *Cognitive Science: A Multidisciplinary Journal*, 7(2): 155-170.
- Gibbs, R.** (2008), *The Cambridge Handbook of Metaphor and Thought*, Cambridge University Press.
- Giere, R** (1999), 'Using Models to Represent Reality', in Ed. L. Magnani, N. J. Nersessian, and P. Thagard (1999).
- (2004), 'How Models are Used to Represent Reality', *Philosophy of Science* 71(5): 742-752
- (forthcoming), 'An Agent-Based Conception of Models and Scientific Representation', *Synthese*.

- Glucksberg, S.** (2008), 'How metaphors create categories – quickly', in Gibbs, R. (2008).
- Gödel, K.** (1947), What is Cantors continuum problem? In P. Benacerraf, & H. Putnam (Eds.), (1984).
- Godfrey-Smith, P.** (2009), 'Models and Fictions in Science', *Philosophical Studies* 143: 101-116.
- Goldman, A.** (1967), 'A Causal Theory of Knowing,' *Journal of Philosophy* 64, pp. 357-372.
- Goodman, N.** (1977), *Structure of Appearance*, 3rd ed. Dordrecht: Reidel.
- Greco, J.** (2010) *Achieving Knowledge*. New York: Cambridge University Press.
- (forthcoming) 'Episteme, Knowledge, and Understanding'.
- Hackl, M. and Thornton, R. (eds.)** (forthcoming), *Asserting, Meaning, and Implying*.
- Haddock, A., Millar, A., Pritchard, D. (eds)** (2009), *Epistemic Value*, OUP.
- Hale, B.** (1994), 'Is Platonism Epistemologically Bankrupt?', *Philosophical Review* 103, pp.299-325.
- Hale, B. & Wright, C.** (2001), *The Reason's Proper Study*. Oxford: OUP.
- (2005), 'Logicism in the Twenty-First Century', in *The Oxford Handbook of Philosophy of Mathematics and Logic*: OUP
- (2009), 'The Metaontology of Abstraction', in *Metametaphysics*, ed. Chalmers, D., Manley, D. and Wasserman, R.: OUP.
- Held, C.** (2007), 'When Does a Scientific Theory Describe Reality?', in Suárez, M. (ed.) (2007).
- Hellman, G.** (1989), *Mathematics without Numbers: Towards a Modal-Structural Interpretation*, Clarendon Press.
- Hempel, C.** (1965), *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*, New York: Free Press.
- Hirsch, E.** (1993), *Dividing Reality*. OUP
- (2009), 'Ontology and Alternative Languages', in *Metametaphysics*, ed. Chalmers, D., Manley, D. and Wasserman, R.: OUP
- Hofweber, T.** (2005), 'A Puzzle about Ontology', *Noûs* 39: 256-83.
- (2007), 'Innocent Statements and their Metaphysically Loaded Counterparts', *Philosophers' Imprint*, 7:1.
- Horwich, P.** (1998), *Truth*, 2nd edition, OUP.
- Howson, C.** (2000), *Humes Problem: Induction and the Justification of Belief*. Oxford: Oxford University Press.
- James, W.** (1907), 'Pragmatism's Conception of Truth', in Blackburn & Simmons (1999).

- Kant, E.** (1781), *Critique of Pure Reason*.
- Kalderon, M.** (2005), *Moral Fictionalism*, OUP.
- Ketland, J.** (1998), *The Mathematization of Nature*. Ph.D. dissertation: LSE.
- (2004), ‘Empirical Adequacy and Ramsification’, *Brit. J. Phil. Sci.* 55: 287-300.
- (2005), ‘Some more curious inferences’, *Analysis*, 65:1, pp.18-24.
- (2011), ‘Nominalistic Adequacy’, *Proceedings of the Aristotelian Society*, 11:2, pp. 201-217.
- Kitcher, P.** (1989), ‘Explanatory Unification and the Causal Structure of the World’, in *Scientific Explanation*, P. Kitcher and W. Salmon, 410505. Minneapolis: University of Minnesota Press.
- (2002). ‘Scientific Knowledge, *Oxford Handbook of Epistemology*, (ed.) P. Moser, Oxford University Press, Oxford.
- Krifka, M. et al.** (1995), ‘Genericity: An Introduction’ in Carlson, G. and Pelletier, F. (eds.) (1995), *The Generic Book*, Chicago: University of Chicago Press.
- Knuuttila, T.** (2007), ‘Representation, Idealization, and Fiction in Economics: From the Assumptions Issue to the Epistemology of Modelling’, in Suárez, M. (ed.) (2007).
- Kvanvig, J.** (2003), *The Value of Knowledge and the Pursuit of Understanding*, Cambridge University Press.
- (2005), ‘Is Truth the Primary Epistemic Goal?’, in Steup, M. and Sosa, E. (2005).
- Lackey, J.** (2008), ‘What Luck is Not’, *Australasian Journal of Philosophy* 86:2, pp. 255–267.
- Lakoff, G.** (1992), ‘The Contemporary Theory of Metaphor’, In Ortony, A. (ed.) *Metaphor and Thought* (Second Edition). Cambridge University Press.
- (ongoing), URL = <<http://cogsci.berkeley.edu/lakoff/MetaphorHome.html>>
- Lakoff, G. & Johnson, M.** (1980a), ‘Conceptual Metaphor in Everyday Language’, *The Journal of Philosophy*, 77: 8, pp.453-486.
- (1980b), *Metaphors We Live By*. Chicago: University of Chicago Press.
- (2000), *Philosophy in the Flesh*. Cambridge, Mass.: MIT Press.
- Ladyman, J. and Ross, D. with Spurrett, D. and Collier, J.** (2007). *Every Thing Must Go*, Clarendon Press.
- Leng, M.** (2002), ‘Whats Wrong with Indispensability? (Or, The Case for Recreational Mathematics)’, *Synthese* 131, pp.395-417.
- (2005), ‘Revolutionary Fictionalism: A Call to Arms’, *Philosophia Mathematica* 3: 13, pp.277-293.

- (2007), ‘What’s there to know? A fictionalist account of mathematical knowledge’, in Leng, M., Paseau, A. and Potter, M. (eds.) (2007).
- (2010), *Mathematics and Reality*, OUP.
- Leng, M., Paseau, A. and Potter, M. (eds.)** (2007), *Mathematical Knowledge*, OUP: New York.
- Leslie, S.** (2007), ‘Generics and the structure of the mind’, *Philosophical Perspectives*, 21(1).
- (2008), ‘Generics: Cognition and acquisition’, *Philosophical Review*, 117(1).
- Leslie, S., Khemlani, S., Prasada, S., and Glucksberg, S.** (2009), ‘Conceptual and linguistic distinctions between singular and plural generics’, *Proceedings of the 31st Annual Cognitive Science Society*. Amsterdam: Cognitive Science Society.
- Lewis, D.** (1973). *Counterfactuals*, Oxford: Blackwell.
- (1979). Counterfactual Dependence and Times Arrow, *Noûs*, 13: 455-76.
- (1986), *On the Plurality of Worlds*, Blackwell, Oxford.
- (1991), *Parts of Classes*. Blackwell, Oxford.
- (1993), ‘Mathematics is Megathology’, *Philosophia Mathematica*, 3:1, pp.3-23.
- Lipton, P.** (2002), ‘Hume’s Problem: Induction and the Justification of Belief’, *BJPS*, 53(4): 579-583.
- (2004), *Inference to the Best Explanation* (Second Edition). Routledge.
- Lin, C. and Segel, L.** (1988), *Mathematics Applied to Deterministic Problems in the Natural Sciences*, SIAM: Philadelphia.
- Luce, R. D., D. H. Krantz, P. Suppes, and A. Tversky** (1971). *Foundations of measurement*. San Diego, CA: Academic Press.
- Maddy, P.** (1997), *Naturalism in Mathematics*, Clarendon Press.
- Magnus, P. D. & Callender, C.** (2004), ‘Realist Ennui and the Base Rate Fallacy’, *Philosophy of Science*, 71, pp.320-338.
- Markosian, N.** (2008), ‘Restricted Composition’, in Sider, T., Hawthorne, J. and Zimmerman, D (eds.) (2008).
- Matthews, R.** (1994) ‘The Measure of the Mind’, *Mind* 103(410): 131-146.
- (2007), *The Measure of the Mind: Propositional Attitudes and their Attribution*, Clarendon Press.
- Maudlin, T.** (2007). *The Metaphysics Within Physics*, Clarendon Press.
- Melia, J.** (1995), ‘On What Theres Not’, *Analysis*, 55: 4, pp.223-229.
- (2000), ‘Weaseling Away the Indispensability Argument’, *Mind*, 109: 435, pp.455-479.

— (2002), 'Response to Colyvan', *Mind*, Vol. 111 . 441, pp.75-79.

Millikan, R. (1984), *Language, Thought and Other Biological Categories*, MIT Press.

— (1989), 'In Defense of Proper Functions', *Philosophy of Science* 56: 288-302.

Miščević, N. (2007), 'Armchair Luck: Apriority, intellection and epistemic luck', *Acta Analytica*, 22:1, pp.48-73.

Moltmann, F. (2003). *Propositional Attitudes Without Propositions*, *Synthese*, 135: 77-118.

— (forthcoming), 'Reference to Numbers in Natural Language', *Philosophical Studies*.

Morgenbesser, S. (1967), *Philosophy of Science Today*, New York: Basic Books.

Morris, T. and Menzel, C. (1986), 'Absolute Creation', *American Philosophical Quarterly*.

Murray, J. et al. (eds.) (1971), *The Oxford English Dictionary*, Oxford: OUP.

Neta, R. and Rohrbaugh, G. (2004), 'Luminosity and the safety of knowledge', *Pacific Philosophical Quarterly* 85 pp.396-406.

Nietzsche, F., Hollingdale, R. (tr.) (1973), *Beyond Good and Evil*, Penguin Books.

Nozick, R. (1981), *Philosophical Explanations*, Cambridge: Cambridge University Press.

Paseau, A. (2005), 'Naturalism in Mathematics and the Authority of Philosophy', *British Journal for the Philosophy of Science* 56, pp.377-396.

Peirce, C.S. (1878), 'How to Make Our Ideas Clear', *Popular Science Monthly* 12, pp.286-302.

Pinker, S. (2007), *The Stuff of Thought: Language as a Window to Human Nature*. Penguin Books.

Pincock, C. (2004), 'A New Perspective on the Problem of Applying Mathematics', *Philosophia Mathematica* 3:12, pp.135-161.

— (2007), 'A Role for Mathematics in the Physical Sciences', *Noûs* 41:2, pp.253-275.

— (2009), 'Towards a Philosophy of Applied Mathematics', in Bueno and Linnebo (eds.) (2009).

Plantinga, A. (1974), *The Nature of Necessity*, Oxford: Oxford University Press.

— (1993), *Warrant: The Current Debate*. New York: Oxford University Press.

— , (2012), *Where the Conflict Really Lies*, OUP: USA

— (unpublished), 'Naturalism Defeated', available on-line at <http://www.homestead.com/philofreligion/files/alspaper.htm>

Pollock, J. (1986), *Contemporary Theories of Knowledge*, Savage, MD: Rowman and Little-Field.

Potter, M. (2007) 'What is the problem of mathematical knowledge?', in *Mathematical Knowledge*, Leng, M, Paseau, A. and Potter, M. (eds.).

- Price, H.** (1997), 'Carnap, Quine and the Fate of Metaphysics', *Electronic Journal of Analytic Philosophy*, 5.
- (2004), 'Naturalism Without Representationalism', in Macarthur, D. and Caro, M. (eds.) (2004).
- (2009), 'Metaphysics after Carnap: The Ghost Who Walks?', in Chalmers, D., Manley, D. and Wasserman, R. (eds.) (2009).
- Priest, G.** (1997), 'Sylvan's Box: A Short Story and Ten Morals', *Notre Dame Journal of Formal Logic*, 38: pp.573-81.
- (2005), *Towards Non-Being: the logic and metaphysics of intentionality*. New York: Oxford University Press.
- (2008), 'The Closing of the Mind: How the Particular Quantifier Became Existentially Loaded Behind Our Backs', *The Review of Symbolic Logic* 1:1, pp. 42-55.
- Pritchard, D.** (2005), *Epistemic Luck*. New York: Oxford University Press
- (2007), 'Anti-Luck Epistemology', *Synthese* 158, pp.277-97
- (2008), 'Safety-Based Epistemology: Whither Now?', *Journal of Philosophical Research* 33.
- (forthcoming), 'Knowledge, Understanding and Epistemic Value', to appear in *Epistemology*.
- Putnam, H.** (1971): 'Philosophy of Logic', in *Mathematics, Matter and Method: Philosophical Papers Volume 1*, 2nd edn, Cambridge: Cambridge University Press, 1979, pp.323-57. (Originally published as *Philosophy of Logic*, New York: Harper Torchbooks, (1971.)
- (1981), *Reason, Truth and History*, Cambridge University Press.
- (1985), 'A Comparison of Something with Something Else', in *Words and Life* (1994), Cambridge, Mass.: Harvard University Press.
- (1987), *The Many Faces of Realism*, Open Court.
- (2004), *Ethics Without Ontology*, Harvard University Press.
- Quine, W.V.O.** (1948), 'On what there is', *Review of Metaphysics* 2, in Quine (1961).
- (1953), *From a Logical Point of View*. Harvard University Press.
- (1960), *Word and Object*, MIT Press.
- (1975), 'Five milestones of empiricism', in Quine (1981), pp.67-72.
- (1978), 'Facts of the Matter', *Southwestern Journal of Philosophy*.
- (1981), *Theories and Things*, Harvard, Cambridge, Massachusetts.
- (1986), *Philosophy of Logic*, 2nd ed. Cambridge, Mass.: Harvard University Press.

— (1990), *Pursuit of Truth*, Harvard University Press.

Radder, Hans (ed.) (2003) *The Philosophy of Scientific Experimentation*, Pittsburgh: University of Pittsburgh Press.

Rayo, A. (2009), 'Towards a Trivialist Account of Mathematics', in *New Waves in the Philosophy of Mathematics*, (2009), Palgrave MacMillan.

Resnik, M. (1997), *Mathematics as a Science of Patterns*, Clarendon Press.

Recanati, F. (2004), *Literal Meaning*, Cambridge University Press.

Reichenbach, H. (1965), *The Theory of Relativity and A Priori Knowledge*, Tr. Maria Reichenbach, Berkeley: University of California Press.

Richard, M. (1990) *Propositional Attitudes: An Essay on Thoughts and How We Ascribe Them*, Cambridge: Cambridge University Press.

Rosen, G. (1990), 'Modal Fictionalism', *Mind*. 99:395, pp. 327-354.

— (2001), 'Nominalism, Naturalism, Epistemic Relativism', *Philosophical Perspectives*, 15, pp.69-91.

— (2012), 'Abstract Objects', *The Stanford Encyclopedia of Philosophy*, Zalta, E. (ed.), URL = <http://plato.stanford.edu/entries/abstract-objects/>.

Rouse, J. (1987), (1987). *Knowledge and power*, Ithaca, NY: Cornell University Press.

— (2007), 'Laboratory Fictions', in Suárez (ed.) (2007).

Russell, B. (1913). 'On the Notion of Cause', *Proceedings of the Aristotelian Society, New Series*, 13: 1-26.

Saatsi, J. (2011), 'The Enhanced Indispensability Argument: Representational versus Explanatory Role of Mathematics in Science.' *BJPS*, 62: 143-154.

Saatsi, J. and French, G. (eds.) , *The Continuum Companion to the Philosophy of Science*, Continuum Publishing Corporation.

Sainsbury, R.M. (1997), 'Easy Possibilities', *Philosophy and Phenomenological Research*, 57:4, pp. 907-919.

Salmon, W. (1984), *Scientific Explanation and the Causal Structure of the World*, Princeton: Princeton University Press.

— (1994) 'Causality Without Counterfactuals', *Philosophy of Science*, 61: 297-312.

— (1997), 'Causality and Explanation: A Reply to Two Critiques', *Philosophy of Science*, 64: 461-477.

Sellars, W. (1953), 'A Semantical Solution to the Mind-Body Problem', *Methodos* 5: 45-84.

— (1963). 'Abstract Entities', *Review of Metaphysics* 16: 627-671.

- Shaffer, J.** (2009), 'On What Grounds What', in Chalmers, D., Manley, D. and Wasserman, R. (eds.) (2009).
- Shapiro, S.** (1983), 'Mathematics and Reality', *Philosophy of Science*, 50: 523-548.
- (1997), *Philosophy of Mathematics: Structure and Ontology*, OUP.
- (2005), *The Oxford Handbook of Philosophy of Mathematics and Logic*, OUP: New York.
- Sidelle, A.** (2002), 'Is there a true metaphysics of material objects?', *Philosophical Issues*, 12: 118-45.
- Sider, T.** (2001), *Four-Dimensionalism*, OUP
- (2009) 'Ontological Realism', in *Metametaphysics*, ed. Chalmers, D., Manley, D. and Wasserman, R.: OUP
- Sider, T., Hawthorne, J. and Zimmerman, D (eds.)** (2008). *Contemporary Debates in Metaphysics*, Blackwell Publishing.
- Smart, J.** (1968), *Between Science and Philosophy*. New York: Random House.
- Soames, S.** , (2003), *Philosophical Analysis in the Twentieth Century: Volume 2, The Age of Meaning*, Princeton University Press.
- (2005), 'Naming and Asserting', in Szabó, Z., (ed.) (2005).
- (2007), 'Actually', *Proceedings of the Aristotelian Society Supplementary Volume*
- (2010), *Philosophy of Language (Princeton Foundations of Contemporary Philosophy)*, Princeton University Press.
- (2010a), *What is Meaning?* (Soochow University Lectures in Philosophy), Princeton University Press.
- (forthcoming), 'The Gap Between Meaning and Assertion: Why what we literally say often differs from what our words literally mean', in Hackl, M. and Thornton, R. (eds.) (forthcoming).
- Sober, E.** (1993), 'Mathematics and Indispensability,' *The Philosophical Review* 102:1, pp.35-57.
- Sosa, E.** (1999), 'How to Defeat Opposition to Moore', *Noûs*, 33:12, pp.141-153.
- Sperber, D. and Wilson, D.** (1986), *Relevance*, Harvard University Press.
- (2008), 'A deflationary account of metaphors', in Gibbs, R. (2008).
- Stalnaker, R.** (1968). *A Theory of Conditionals*, *Studies in Logical Theory: American Philosophical Quarterly Monograph Series*, No. 2. Oxford: Blackwell.
- (1970), 'Pragmatics', *Synthese* 22: 272-289.
- (1976), 'Possible Worlds', *Noûs* 10, pp. 65-75

— (1984) *Inquiry*, Cambridge, MA: MIT Press.

Steup, M. and Sosa, E. (eds.) (2005), *Contemporary Debates in Epistemology*, Blackwell.

Suárez, M. (1999a) 'The Role of Models in the Application of Scientific Theories: Epistemological Implications' in *Models as Mediating Instruments*, M. Morgan and M. Morrison (eds.), pp. 168-196, Cambridge University Press.

— (1999b) 'Theories, Models and Representations', in *Model-Based Reasoning in Scientific Discovery*, L. Magnani, N. Nersessian and P. Thagard (eds), pp. 75-83, Plenum, New York.

— (2003) 'Scientific Representation: Against Similarity and Isomorphism', *International Studies in the Philosophy of Science*, 17, 3, pp. 225-244.

— (2004) 'An Inferential Conception of Scientific Representation', *Philosophy of Science*, 71, pp. 767-779.

— (2007), (ed.) *Fictions in Science: Philosophical Essays on Modeling and Idealization*, Routledge.

— (2008) 'Theories: Tools vs. Models' (with Nancy Cartwright), *Studies in the History and Philosophy of Modern Physics*, 39, pp. 62-81.

— (2009a) 'Fictions in Scientific Practice', in *Fictions in Science: Philosophical Essays on Modeling and Idealisation*, Routledge, pp. 1-15.

— (2009b) 'Scientific Fictions as Rules of Inference', in M. Suárez (ed.), *Fictions in Science: Philosophical Essays on Modeling and Idealization*, Routledge, pp. 158-178.

— (2010a) 'Scientific Representation', *Philosophy Compass* 5 (1): pp. 91-101. Blackwell.

— (2010b), 'Fictions, Inference, and Realism', in J. Woods (ed.), *Fictions and Models: New Essays*, Philosophia Verlag, Munich, pp. 225-245.

Suppes, P. (1960), 'A Comparison of the Meaning and Uses of Models in Mathematics and the Empirical Sciences', *Synthese* 12: 287-301.

— (1967), 'What is a Scientific Theory?', in Morgenbesser (1967).

— (2003), *Representation and Invariance of Scientific Structures*, University of Chicago Press.

Suppes, P. and Zinnes, J. (1963) 'Basic Measurement Theory', in R. Luce, D. Krantz, and E. Galanter (eds.) *Handbook of Mathematical Psychology*, Vol. I (New York: Wiley and Sons), 1-76.

Syngé, J. L. & Griffith, B. A. (1942), *Principles of Mechanics*. (McGraw-Hill Book Company Inc., New York, Third edition, 1959.

Szabó, Z., (ed.) (2005), *Semantics vs. Pragmatics*, OUP: New York.

Teller, P. (2001), 'Twilight of the Perfect Model Model', *Erkenntnis*.

- Thomasson, A.** (2007), *Ordinary Objects*. OUP
- (2009), Answerable and Unanswerable Questions, in *Metametaphysics*, ed. Chalmers, D., Manley, D. and Wasserman, R.: OUP
- van Cleve, J.** (2008) ‘The Moon and Sixpence: A Defense of Mereological Universalism’. In Sider, T., Hawthorne, J. and Zimmerman, D (eds.), (2008).
- van Fraassen, B.** (1980), *The Scientific Image*, Clarendon Press.
- (2008), *Scientific Representation: Paradoxes of Perspective*, OUP.
- van Inwagen, P.** (1990), *Material Beings*. Ithaca, NY: Cornell University Press
- (1998), ‘Meta-ontology’, *Erkenntnis* 48: 233-50.
- (2000), ‘Quantification and Fictional Discourse’, in Everet, A. and Hofweber, T. (2000).
- (2009), ‘Being, Existence and Ontological Commitment’, in *Metametaphysics*, ed. Chalmers, D., Manley, D. and Wasserman, R.: OUP
- Waterfield, R. (tr.)** (2005), *Plato: Meno and Other Dialogues*, OUP.
- Weatherson, B.** (2004), ‘Luminous Margins’, *Australasian Journal of Philosophy*, 82:3, pp.373-383.
- Wetzel, L.** , ‘The Trouble With Nominalism’, *Philosophical Studies*, 98:3, pp.361-370.
- (2006), ‘Types and Tokens’, in *The Stanford Encyclopedia of Philosophy*, Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/entries/types-tokens/>>
- (2009), *Types and Tokens*. MIT Press
- Williamson, T.** (2000), *Knowledge and its Limits*, Oxford: Oxford University Press.
- Wisnovsky, R.** (2005), ‘Avicenna and the Avicennian Tradition’, in Adamson P. and Taylor, R. (eds) (2005).
- Woods, J. (ed.)** (2010), *Fictions and Models: New Essays*, Munich: Philosophia Verlag.
- Woodward, J.** (2003), *Making Things Happen: A Theory of Causal Explanation*, OUP: Oxford.
- Yablo, S.** (1998), ‘Does Ontology Rest on a Mistake?’, *Aristotelian Society Supplementary Volume*, 72: 229-262.
- (2009), ‘Must existence-questions have answers?’, in *Metametaphysics*, ed. Chalmers, D., Manley, D. and Wasserman, R.: OUP
- Zagzebski, L.** (1996), *Virtues of the Mind*, Cambridge: Cambridge University Press