

STRUCTURING REALITY

By

NAOMI MARGARET CLAIRE THOMPSON

A thesis submitted to the University of Birmingham for the degree of
DOCTOR OF PHILOSOPHY

Department of Philosophy
School of Philosophy, Theology and Religion
College of Arts and Law
University of Birmingham
July 2014

Abstract

This thesis explores attempts to characterise the structure of reality. Three notions stand out: Lewisian *naturalness*, Sider's '*structure*', and *grounding*, where the latter has become the most popular way to characterise the structure of reality in the contemporary literature. I argue that none of these notions, as they are currently understood, are suited for limning the metaphysical structure of reality. In the first part of the thesis I argue that, by the lights of the relevant theories, both naturalness and structure fall short of the theoretical role carved out for those posits. In the second part of the thesis I present two challenges to the 'orthodox' conception of grounding. The first contests the standard assumption that grounding is asymmetric, both by citing what I take to be best described as symmetric instances of grounding, and by developing and arguing for a new theory of metaphysical structure – 'metaphysical interdependence' – which takes grounding to be nonsymmetric. The second challenge concerns the relationship between grounding and (metaphysical) explanation, and leads to a dilemma for the grounding theorist. My proposed resolution to the dilemma is to adopt an antirealist approach to grounding, which I further motivate and develop in the final chapter.

Acknowledgements

I would like first to acknowledge the support of three supervisors I have had the pleasure to work with over the course of writing this thesis: Yujin Nagasawa, Nikk Effingham, and, in particular, Alastair Wilson, whose dedication to reading and discussing every aspect of this thesis (sometimes numerous times) have helped me immensely. I would also like to thank Nicholas Jones for providing me with extremely helpful comments and suggestions on what have become large sections of my thesis, and Kirk Surgener, who proofread the entire thesis at the last minute, and offered numerous philosophical suggestions. I would also like to thank staff and students at the University of Birmingham for providing a friendly, supportive, and stimulating environment in which to conduct my research, and participants and organisers at the many conferences and workshops at which I have presented aspects of it.

For their unfailing love and support I would like to thank my family, and for their support and encouragement I would like to thank my friends. Most of all I wish to thank Darragh Byrne, who not only helped a great deal with the philosophical content of this thesis, but also provided love, support, and encouragement, and who made the last few months a great deal easier than they would have been otherwise.

This research was made possible through funding from the Arts and Humanities Research Council.

Table of Contents

Introduction.....	1
Chapter 1: Preliminaries.....	4
1. Ontological dependence	4
1.1. Existential dependence.....	5
1.2. Essential dependence.....	6
2. Grounding.....	8
2.1. Regimenting grounding claims.....	9
2.2. Responding to scepticism about grounding.....	15
2.3. Connections.....	23
3. Fundamentality.....	29
3.1. Ungroundedness.....	29
3.2. Ontological commitment	33
3.3. Primitivness	34
3.4. Linguistic conception.....	35
3.5. Naturalness/structure.....	36
3.6. Truth-makers.....	38
3.7. Explanation	39
4. Concluding Remarks	40
Chapter 2: Naturalness and Structure.....	41
1. Naturalness.....	41
2. The naturalness role.....	43
2.1. Empiricism and laws.....	45
2.2. Supervenience	46
2.3. Intrinsicity and Duplication.....	51
2.4. Reference magnetism.....	53
3. How natural is perfect naturalness?	56

4.	Consequences	60
5.	Structure	62
6.	Conclusion	67
Chapter 3: Metaphysical Interdependence.....		68
1.	Foundationalism, infinitism, interdependence	70
2.	Examples of nonsymmetric grounding	75
3.	Arguments for metaphysical interdependence.....	79
3.1.	Part-whole grounding.....	79
3.2.	Gunk and junk.....	80
3.3.	Theoretical virtue.....	82
3.4.	Grounding grounding.....	88
4.	Conclusion	90
Chapter 4: Metaphysical Explanation		92
1.	Grounding and explanation.....	92
2.	Models of explanation.....	97
2.1.	The covering-law model.....	98
2.2.	The causal model.....	103
2.3.	The pragmatic approach.....	106
2.4.	Metaphysical explanation and explanation-seekers.....	109
3.	Articulating the problem.....	111
4.	First horn: metaphysical explanation	112
5.	Second horn: the grounding-explanation connection.....	118
6.	Conclusion	124
Chapter 5: Antirealism about Grounding		125
1.	Attitudes to grounding	125
2.	Motivations for antirealism	127
2.1.	Intuition	128
2.2.	Explanation	131

2.3.	Antirealism or eliminativism?	136
3.	Types of antirealism.....	138
3.1.	Subjectivism.....	138
3.2.	Non-cognitivism.....	140
3.3.	Error theory.....	142
3.4.	Fictionalism	142
4.	Expressivism about grounding	143
4.1.	Credence model.....	144
4.2.	Intelligibility model.....	149
4.3.	The Frege-Geach problem.....	151
5.	Fictionalism.....	155
5.1.	Field-style fictionalism about grounding	156
5.2.	Yablo-style fictionalism about grounding.....	160
6.	Comparative remarks	164
7.	Realism, sophisticated antirealism, and grounding.....	165
7.1.	The problem of ‘creeping minimalism’.....	166
7.2.	Fine’s ‘Question of Realism’	169
7.3.	Sophisticated Antirealism about grounding	175
8.	Conclusion	179
	Concluding Remarks: Antirealism, Nonsymmetry, Explanation and Structure	180
	Appendix: Reference Magnetism	183
1.	Putnam’s paradox.....	183
2.	Reference magnets and fundamental physics.....	187
3.	Use and Eligibility	189
	References	192

Introduction

This thesis explores some of the notions that have been appealed to in order to characterise the structure of reality. The most significant of these are the notions of *naturalness* and of *grounding*. I argue that naturalness theory, and an influential extension of naturalness theory defended by Ted Sider are both ill-suited to play the role carved out for them – that of limning the structure of reality. In the second, and main part of the thesis I move on to discuss grounding, where I argue against the ‘orthodox’ conception of that notion. In the final chapter I present a new and somewhat radical positive proposal for understanding grounding in an antirealist way.

The first chapter begins with a discussion of the notion of ontological dependence, distinguishing between existential and essential dependence, and summarising a number of different accounts thereof. This paves the way for introduction of the notion of grounding, in terms of which ontological dependence has come to be most often characterised. The second part of Chapter 1 is given over to exploring various aspects of the notion of ground, including how such claims are to be formulated and regimented. Some philosophers are sceptical about grounding, and so I introduce the main sceptical arguments that are found in the literature, and outline some response-strategies. I then explore connections between grounding and other metaphysical notions: modality; essence; truth-making; and reduction. In the final part of the first chapter I discuss a number of different ways to characterise the notion of fundamentality with which ontological dependence and grounding are very often associated.

The focus of the second chapter is naturalness. David Lewis carves out a theoretical role for an elite class of properties that he terms the ‘perfectly natural’ properties, and which he takes to be the properties that do the most important work in metaphysics (and science). I explain some of central tenets of naturalness theory, and I argue that by the lights of naturalness theory, the property *being a perfectly natural property* ought itself to be perfectly natural, but that also by the lights of naturalness theory, it cannot be perfectly natural (or, in fact, even quite natural). This renders naturalness unsuitable for describing

reality's structure. I argue that Sider's extension of naturalness, discussed in terms of his theoretical posit 'structure' is unsuitable for somewhat similar reasons. The notion of structure that Sider takes to be primitive is that of *absolute* structure, but it is a notion of relative structure that Sider takes to do most of the interesting work. Absolute structure is not absolutely structural, and so Sider's extension of Lewisian naturalness is subject to an analogous objection.

In Chapter 3, the focus shifts to grounding. I argue against the orthodox view that grounding is an asymmetric relation, and develop an account of the metaphysical structure of the world that results from dropping the asymmetry constraint on grounding. I claim that this account has some benefits over the traditional models, and so the idea that grounding is a nonsymmetric relation ought to be taken seriously. The antirealist approaches to grounding I develop in the final chapter are consistent with (though not reliant upon) this nonsymmetric notion of ground.

Chapter 4 explores connections between grounding and (metaphysical) explanation. It is very common to find appeals in the literature to the alleged explanatory character of ground, but in this chapter I highlight a tension between that connection, the view that grounding is a feature of mind-independent reality, and the pragmatic features of the relevant sort of explanation. That the relevant sort of explanation has pragmatic features is a claim I argue for in some detail. I conclude that neither denying the link between grounding and metaphysical explanation, nor denying that metaphysical explanation has pragmatic features can be made attractive. This tension then forms part of my argument for antirealism about grounding, because the problem is resolved if grounding is not a feature of mind independent reality.

In the final chapter I first offer some additional motivations for antirealism about grounding, before considering which varieties of antirealism might be most appropriate for the grounding case. I then develop two models for expressivism about grounding, and two models for fictionalism about grounding. I explain how each of these models is able to account for various features we take grounding to have, and how defenders of each model might respond to some apparent objections. I then focus in detail on one theoretical role grounding has been taken to play – the role of distinguishing realism from sophisticated antirealism – and argue that grounding as understood in accordance with an antirealist

approach is still well suited to play that role. Finally, I explain how antirealists about grounding are able to accommodate the claims about nonsymmetry and explanation that I made in Chapters 3 and 4.

I conclude that a notion of grounding that takes metaphysical structure to be a projection of our cognitive architecture onto the world is capable of playing the roles that have led to philosophers positing a notion of grounding, but that it avoids the problematic metaphysical and epistemic commitments of realism about grounding. It is our epistemic and conceptual commitments that paint a structure onto reality, and we communicate that structure using grounding locutions.

Chapter 1: Preliminaries

My intention in this chapter is to map out the key notions that will be the subject of this thesis, and thus to clarify how I will understand them in the rest of the discussion. I begin with a discussion of ontological dependence, before introducing the notion of grounding. I discuss ways to regiment grounding, including outlining the debate over how grounding claims are to be formulated, and about the structural properties that grounding is usually taken to have. I outline some sceptical challenges to the concept of grounding, discuss responses, and assess connections grounding has to related notions. In the final section, I identify and critically discuss a number of different accounts of fundamentality.

1. Ontological dependence

An idea, seeming to originate with Aristotle (e.g. 1998) has recently become the subject of a flurry of interest amongst metaphysicians and those working in related disciplines. The idea is that there is a distinctively metaphysical way one thing might be said to depend on another, i.e. a way in which one thing might depend on another for its existence and nature (cf. Schaffer, 2010a). The notion at stake here is often taken to be primitive, in the sense that it resists reductive analysis. It is usually introduced by example (see e.g. Fine, 1995a). It is the way in which non-empty sets are said to depend on their members, holes to depend on their hosts; mental facts to depend on physical facts; or a smile to depend on the mouth that is smiling. Questions about whether a whole depends on the parts that compose it, or the parts depend on the whole they are part of (see Schaffer, 2010a) are questions about ontological dependence.

I begin with a (simplified) account of the key notions of ontological dependence, in order to trace the recent history of the grounding relation (in terms of which ideas about ontological dependence are now most often expressed).

1.1. Existential dependence

The notion of existential dependence is that of one thing depending on another for its very existence. Following Lowe (2010) and Correia (2008: 1014-1016), we can identify two main ways to define existential dependence.

(1) x depends for its existence on y =_{df} Necessarily, x exists only if y exists.

Note that the notion of dependence in (1) is rigid; the existence of y is strictly implied by the existence of x . This can be contrasted with a non-rigid definition of existential dependence:

(2) x depends for its existence on the F =_{df} Necessarily, x exists only if the F exists.

The non-rigid (or generic) construal of existential dependence allows for cases where the dependent entity might depend on different entities at different possible worlds. For example, the set of people living in France might be said always to existentially depend on its members, but the members of that set might be different at different possible worlds (or different at different times at the same world). Whilst the rigid characterisation of existential dependence requires the existence of a specific object, the non-rigid characterisation requires the existence of an object of a certain sort (Correia, 2008: 1015). Since they make use of a notion of necessity, these accounts of dependence are modalised. Exactly what the accounts amount to is sensitive to the interpretation of the necessity operator. In most discussions of ontological dependence, the relevant account of necessity is metaphysical, though Fine (1995a: 270) suggests that different accounts of dependence are generated by varying the interpretation of the necessity operator (generating notions of, for example, normative dependence and conceptual dependence).

Both Aristotle and Husserl appear to give an account of dependence in existential terms (on Aristotle see Corkum, 2008, and on Husserl see Correia, 2004). A comprehensive discussion of existential dependence can be found in Part III of Simons (1987). In 1995, Kit Fine proclaimed that the modalised existential account of ontological dependence had ‘taken over’ (1995a: 271), and that it was rare to come across an alternative. He argued, however, that this account of ontological dependence is inadequate for picking out the notion we intend to capture when we make dependence claims.

Fine (e.g. 1995a: 171-2) offers some examples to illustrate that the existential account of ontological dependence can be too coarse grained for the required purpose. Consider the set whose sole member is Socrates (Socrates' singleton set). Necessarily, when Socrates exists, so does his singleton set. It follows from this that, necessarily, Socrates exists only if the set exists. According to (1), Socrates depends for his existence on his singleton set (and vice versa). But this, intuitively, is the wrong result. The set with Socrates as its sole member depends on the existence of Socrates, but the converse is not the case. The existence of the member does not depend on the existence of the set.

Further difficulties are encountered in the case of necessary existents, where (1) seems to be trivially satisfied even though we have no reason to suspect any dependence relation between the entities in question. On a modal existentialist construal of dependence it is always the case that, for example, the number 2 exists if Socrates does. Modifications of the account at hand are not promising – we cannot, for example, simply ban necessary existents from being dependees, because there are plausible cases of ontological dependence between necessary existents (e.g. the set which has the number two as its sole member, and the number two itself) (Fine, 1995a: 172).

More generally, it seems that two philosophers could agree on all of the modal facts, but disagree on the dependence relations (Fine 1995a: 172). Fine's example is of philosophers who both think that persons and minds are distinct and that the one could not exist without the other, but whilst one thinks of minds as abstractions from persons and thus dependent on persons, the other thinks of persons as embodied minds and thus dependent on minds. The modalised existentialist account of ontological dependence does not have the resources to make sense of this disagreement.

1.2. Essential dependence

Fine's solution is to propose an alternative account of ontological dependence, which focuses on the *essence* of the dependent object – what it takes for it (the dependent object) to be that object rather than something else. One way to cash out the notion of essential dependence is with the relation Lowe (2010) calls identity dependence, and defines as follows:

- (3) x depends for its identity upon $y =_{df}$ There is a function f such that it is part of the essence of x that x be the f of y .

That the set with Socrates as its sole member depends for its identity on Socrates means that there is some function (being a unit set) such that it is part of the essence of Socrates' singleton that it be the unit set of Socrates (see Lowe, 2010). So long as we can maintain that it is not part of the essence of Socrates that he be the sole member of $\{\text{Socrates}\}$, it seems that Fine's counterexamples can be avoided.

An alternative construal of essential dependence is what Lowe (2010) calls essential existential dependence, and which he defines as follows:

- (4) x depends for its existence on $y =_{df}$ It is part of the essence of x that x exists only if y exists.

This account has it that if Socrates' singleton depends for its existence on Socrates, then it is part of the essence of Socrates' singleton that it exists only if Socrates exists.

A proper understanding of (3) and (4) requires an account of essence. On what Correia (2008: 1017) calls the widely held reductionist conception of essence, what it is for an object to be essentially φ is that if the object exists, then it φ s. The reductionist account entails two conditionals:

- (α) If x is essentially φ then, necessarily, if x exists then it φ s
(β) If, necessarily, x φ s when it exists, then x is essentially φ

Fine (2004) offers counterexamples to the second conditional. For example, it is necessarily the case that Socrates exists when he exists, but it is not the case that Socrates essentially exists. Fine uses this and similar counterexamples to motivate a rejection of the reductionist construal of essence, positing instead what we might consider a definitional account of essence. The basic idea behind this account of essence is that it is true *in virtue of* the identity of x that it φ 's

If we are to retain a distinction between the essentialist existential account of ontological dependence and the existential account of ontological dependence, we require an account of essence not tied to modality. On pain of counterexamples already discussed, it cannot be the case that what it means for x to have a certain property essentially is for it to be necessary that x exists only if it has that property (Fine, 1995a: 273). Fine's alternative account of dependence makes use of the locution 'in virtue of', which

suggests an explanatory form of ontological dependence that will be the subject of the latter (and main) part of this thesis. I introduce this notion of ontological dependence – called ‘grounding’ – in the following section.

2. Grounding

Ontological dependence is now most often discussed in terms of a relation of grounding, where we can express the idea that, for example, sets ontologically depend on their members by stating that sets are *grounded* in their members. The precise nature of the connection between grounding and other notions of ontological dependence is somewhat fraught. Grounding is sometimes taken to be a relation holding between entities of various ontological categories (e.g. Bennett, 2011; 2014 (Bennett talks about ‘building’ relations rather than grounding relations); Cameron, 2008; Schaffer, 2009; 2010a) and sometimes only between facts or propositions (e.g. Fine, 2001; Audi, 2012a; 2012b). In what follows I stay neutral on what are the relata of the grounding relation, as my interest is in aspects of grounding that are common to both accounts.¹ Some take grounding to be one amongst many different relations of ontological dependence (e.g. composition, supervenience, modal entailment). Others are keen to analyse relations like composition in terms of grounding (e.g. Morganti, 2009; Schaffer, 2010a). My primary focus in this thesis is on grounding, and I remain neutral on the way in which grounding relates to ontological dependence.

The notion of grounding is generally taken to be primitive and unanalysable. This opens grounding up to sceptical attack (e.g. Hofweber, 2009; Daly, 2012; J. Wilson, forthcoming), and also motivates defences of grounding talk that focus on the need for a relation of grounding in our metaphysical arsenal (e.g. Schaffer, 2009; Rosen, 2010; Audi, 2012a; 2012b). Some details of both attack and defence are given in section 2.2. The apparent need to take grounding as primitive means it is often introduced by example. Grounding claims can be expressed with a variety of different locutions. Here are some examples of different types of grounding claims (which may or may not be true):

¹ When I make claims that affect the different accounts in different ways, I try to make clear these differences.

- (a) The set {Socrates} is grounded in Socrates himself
- (b) The moral properties of an action are due to its non-moral properties
- (c) Snow's being white makes true the proposition <snow is white>
- (d) The painting is beautiful in virtue of facts about the way it is received
- (e) Legal facts depend on social facts
- (f) The shape is a square because it has four sides of equal length

Claims such as these all seem to contain an explanatory element (e.g. it is the way that the painting is received that *explains why* it is beautiful) and they are not plausibly explained in causal terms (Socrates' existence and nature does not cause that of the set {Socrates}).² Locutions like 'because' and 'in virtue of' are often encountered in ordinary language as well as in philosophy, and so grounding is thought to be a semi-technical notion that is sometimes expressed in folk discourse (e.g. 'it's wrong to hit your brother because you might hurt him').

2.1. Regimenting grounding claims

There are two main ways in which grounding claims are expressed (see Clark and Liggins, 2012: 816). The first is the predicate theory, where grounding is expressed with a relational predicate such as 'grounds' or 'depends on'. Friends of this approach include Schaffer (2009; 2010a) and Rosen (2010). The second is the connective theory (also called the operational view – see Correia and Schnieder, 2012: 11), where grounding is expressed with a sentential connective like 'because' or 'in virtue of'. This locution is preferred by Fine (e.g. 2001; 2012a) and deRosset (2013). For my purposes here, it will not usually be necessary to take a stand on whether grounding is expressed using a predicate or a sentential connective, and so my inclination is once again to remain neutral as far as possible. Some of what I say is more easily discussed if grounding is treated as a predicate, and so it might seem as though that is my preferred formulation. I take it though that what I say about grounding expressed using a predicate can be reinterpreted taking grounding to be a sentential connective (and vice versa). In some cases (where I think it is required) I make this reinterpretation explicit.

² A. Wilson (manuscript) is an exception to the generally accepted view that grounding is a non-causal relation.

I understand grounding to be a one-many relation, so that f might be grounded in g , or in g, g' , and g'' . When I make a claim of the form g grounds f , I do not mean to suppose that g is the sole ground of f . Unless otherwise stated, it is left open that g only partially ground f . When some set of entities taken together is sufficient to ground some entity x , I call that set of entities a complete ground for x . This distinction maps on to a distinction made in Fine (e.g. 2012a: 50) between full and partial ground. In Fine's system, A is a partial ground for C if A , on its own or with some other truths, is a full ground of C . A full ground for C is sufficient, by itself, to ground the truth of C . Fine argues that the notion of full ground is more basic than that of partial ground, because whilst the partial grounds of $A \wedge B$ and $A \vee B$ are the same (both A and B when both A and B are the case), the full grounds are different (each of A and B is a full ground for $A \vee B$, but the full ground for $A \wedge B$ is both A and B). If we only appeal to partial grounds, we will be unable to distinguish between the full grounds of $A \wedge B$ and $A \vee B$ (Fine, 2012a: 50), and so it is sensible to derive the notion of a partial ground from that of a full ground.

Trogon (2013a) argues for the seemingly common assumption that grounding is a necessary relation, in the sense that propositions corresponding to full grounds modally entail the propositions that they ground (Trogon, 2013a: 466). (If grounding relations obtain between entities, the orthodox view is that the existence of the entities that jointly fully ground some other entity necessitate the existence of that other entity). Trogon's argument depends on a connection between grounding and essence. Since this is the topic of section 2.3.2, I do not discuss it further here.

The orthodox view about the structural principles governing the grounding relation is that it is transitive, irreflexive and asymmetric (see, for example, Schaffer, 2009 and Rosen, 2010). We can define these principles as follows:

Transitivity: If a grounds b , and b grounds c , then a grounds c

Irreflexivity: No entity grounds itself

Asymmetry: If a grounds b (and $a \neq b$) then b does not ground a

If grounding is transitive, irreflexive and asymmetric, then grounding is a strict partial ordering on entities (see section 3.1 below, and Chapter 3 for some relevant discussion). Below I briefly examine the

principles of transitivity and irreflexivity, and discuss some attempts to call them into question. A rejection of the asymmetry of grounding is the subject of Chapter 3 of this thesis, and so I do not discuss it here.

2.1.1. Transitivity

In many cases (e.g. Schaffer, 2009; Fine, 2010) transitivity has been assumed without argument. Indeed, it is a natural assumption that if, for example, the fact that the table exists is grounded in facts about the existence of the particles that compose it (and their arrangement), and facts about those particles (and their arrangement) are grounded in facts about the existence and arrangement of the subatomic particles that compose the particles, then the fact that the table exists is grounded in facts about the existence (and arrangement) of the subatomic particles.

Schaffer (2012) offers three counterexamples to the transitivity of grounding. The most interesting of these (2012: 126) concerns what we might call a near-sphere – a sphere with a minor dent. The fact that the near-sphere is roughly spherical is (presumably) grounded in the fact that it has maximally determinate shape S. The fact that the near-sphere has shape S is grounded in the fact that it has a dent. By the transitivity of grounding, the fact that the near-sphere has a dent grounds the fact that it is roughly spherical. But this is implausible – the fact that the near-sphere has a dent does not seem to provide any kind of explanation for why it is roughly spherical; it is roughly spherical *despite* being dented.

It seems to me that Schaffer's counterexample capitalises on the understanding that it can be reasonable to cite partial grounds for some fact, relative to some background assumptions. Just as when we give an explanation for some phenomena, we are not required to give a full explanation (consider Lewis' (1986a) view of explanation, whereby a full explanation for some event is nothing less than a complete causal history of that event), some or other salient grounding fact will generally do. Thus, whilst the dent in the near-sphere is relevant to grounding the fact that it has shape S, so is the fact that without the dent, the near-sphere would be a sphere. With this in mind, the fact that the near-sphere has a dent *does* seem to ground the fact that it is roughly spherical (it's just that that fact is a partial rather than a full ground, and is not the partial ground that the context of the question leads us to expect). The scenario is set up so

that what we are really looking for is a ground for the fact that the near-sphere is roughly spherical *rather than spherical*.

The treatment offered above comes close to Schaffer's own suggestion for dealing with the counterexamples to transitivity, which is to adopt a contrastive treatment of grounding. The idea is that we should take the structural principles of grounding to hold between *differences* rather than between facts (Schaffer, 2012: 132). It is not the fact that Φ that grounds the fact that Ψ , but the fact that Φ rather than Φ^* that grounds the fact that Ψ rather than Ψ^* . What the relevant contrast is is fixed by context. The proposal I will argue for towards the end of this thesis comes quite close to Schaffer's own, but the two differ in their accounts of how much pragmatic material should be built into the semantics of 'grounds'. I'll save discussion of my own proposal for Chapters 4 and 5, and end this subsection on the promissory note that there is no need to move to a contrastive treatment of grounding if that proposal is adopted.

2.1.2. Irreflexivity

It is usually assumed that grounding is irreflexive. Rosen (2010: 115) states 'just as no fact can make itself obtain, no fact can play along with other facts in making itself obtain'. This strong notion of irreflexivity rules out any entity being even a partial ground for itself. The irreflexivity of grounding has recently been challenged by Jenkins (2011), but it has previously been suggested (e.g. Lowe, 2010) that we ought to take grounding to be an antisymmetric relation, so as not to rule out cases where some entity depends on itself. Note that irreflexivity is entailed by a strong version of asymmetry – one that omits my qualification above that a and b be non-identical. If a cannot ground b when b grounds a even when a and b are the same entity, then no entity can be a partial ground for itself. On the other hand, if (as I argue in Chapter 3) grounding is nonsymmetric, then it must also be nonreflexive.

Jenkins (2011: 270) argues that whilst utterances of the form ' x grounds x ' (or ' x because x ') sound bad (and it is this sounding-bad that accounts for the assumption of irreflexivity) it is in fact not the case that grounding is irreflexive. Her main argument is that there might be all-true triads like the following (Jenkins, 2011: 270):

- (a) S's pain depends on S's brain state B

(b) S's pain does not depend on S's pain

(c) S's brain state B is identical to S's pain

If dependence (and grounding) is irreflexive, then the above is inconsistent.³ One suggestion Jenkins (2011: 271) makes to deal with this is to claim that grounding is really a four-place relation, holding between a state of affairs, another (possibly identical) state of affairs, a feature or aspect of the first state of affairs and a feature or aspect of the second state of affairs. In this way, it is possible to generate a grounding claim between different aspects of the same state of affairs, and thus to accommodate scenarios such as that in (a) to (c) above. Irreflexivity and reflexivity are properties of binary relations, and so if grounding is a four-place relation then it is not appropriate to label it irreflexive (or indeed reflexive).

Nevertheless, there seems to be a natural generalisation of irreflexivity to a four-place relation. A four-place relation is irreflexive when it cannot be the case that all four argument places are satisfied by the same entity. If we adopt this generalisation, then even if grounding should be considered a four-place relation, Jenkins' purported counterexample does not call the irreflexivity of grounding into question.

An alternative suggestion considered by Jenkins is to claim that grounding is hyperintensional, meaning that grounding creates contexts into which we might not be able to substitute necessarily co-extensive terms *salva veritate*. The hyperintensionality of grounding accounts for why (a) to (c) above might all be true, because even though S's brain state B and her pain are necessarily co-extensional, they might have different grounds. Jenkins does not appeal to the hyperintensionality of grounding to deal with her purportedly problematic triad because, she argues, hyperintensionality does not allow us to make sense of how a single dependence relation can both hold and not hold between the same two things (2011: 271).

The natural response for the friend of irreflexivity to make here is emphasise that it is because grounding is hyperintensional that there can be a grounding relation obtaining between S's brain state and her pain, but to ask how that relation can both obtain and not obtain between the same two things is

³ Perhaps nobody would be willing to accept all of (a) to (c), at least without building in a number of qualifications, but for argument's sake I'll grant Jenkins' example.

to ignore the hyperintensional character of grounding. The relation that does obtain between S's brain state and her pain is *different* to the relation that fails to obtain between S's pain and her pain. My view is that grounding is not irreflexive – I take grounding to be nonreflexive – but I take nonreflexivity to fall out of the nonsymmetry of grounding, for which I argue in Chapter 3.

That grounding is hyperintensional is a point of consensus amongst most recent advocates of grounding. Borrowing a further example from Correia and Schneider (2012: 14), whilst it is plausible to suppose that the fact that {Socrates} exists is grounded in the fact that Socrates exists, it is considered implausible that the fact that {Socrates} exists or some tree is not a tree is grounded in the fact that Socrates exists and every tree is a tree. The hyperintensionality of grounding seems to be a species of the hyperintensionality of explanation. It is plausible to suppose that the window broke because the ball was thrown at it, but perhaps not that the window broke and $2 + 2 = 4$ because the ball was thrown at it or $2 + 2 = 5$.

A further point of consensus about grounding is that it is non-monotonic. If Socrates grounds {Socrates} (or the fact that Socrates exists grounds the fact that {Socrates} exists) it does not follow that Socrates and my big toe ground {Socrates} (or that the fact that Socrates exists and the fact that my big toe exists ground the fact that {Socrates} exists). Like hyperintensionality, non-monotonicity seems to be a feature of explanatory notions in general. If my partner's having an affair explains my anger, it doesn't follow that my partner's having an affair and its being sunny outside explain my anger.

A point of disunity between proponents of grounding concerns whether or not grounding is a *well-founded* ordering. If it is, then there is some set of entities that are not themselves grounded. Connections between grounding and fundamentality have been explored in much of the recent literature on grounding, where a common assumption (e.g. Schaffer, 2009: 373) is that an entity is fundamental when nothing grounds it (see section 3 for a discussion of this and other theories of fundamentality). A relative notion of fundamentality can be generated by appeal to the chain of grounding that connects grounded entities to ungrounded entities, and used to ascribe metaphysical priority. When x is more fundamental than y , x is metaphysically prior to y . Some (e.g. Cameron, 2008) hold that whatever is metaphysically

prior is more real than the derivative, posterior entities. I discuss the different approaches to metaphysical structure informed by these sorts of issues in Chapter 3.

2.2. Responding to scepticism about grounding

As mentioned above, not everybody agrees that it is theoretically important to talk in terms of grounding, or even that there is any such relation as grounding. Clark and Liggins (2012: 814) point out that this resistance towards grounding is more often encountered in conversation than in print, but there are nevertheless some extant dissenting works. Scepticism about grounding (as it is called in the literature) can proceed by different strategies. Daly (2012) argues that the best of these is to claim that grounding talk is unintelligible. J. Wilson (forthcoming) argues that purported instances of grounding are in each case merely an instance of one or other more familiar notion. A similar line is taken by Hofweber (2009), who takes grounding to be an ‘esoteric’ notion, and talk of grounding and related concepts to be ‘absurd’. I discuss these positions in turn in the rest of this section. Other possible attacks that might be made on grounding include claiming that grounding talk is intelligible, but that there are no actual instances of grounding. Alternatively, one might argue that concerns about the epistemology of grounding outweigh any benefits that might be gained by taking grounding to be a worldly relation, but that grounding talk is nevertheless valuable. This latter strategy motivates much of my work in the latter part of this thesis, and so I postpone discussion of it until then.

2.2.1. Intelligibility

Daly (2012: 89) implements a two-stage strategy in his sceptical argument against grounding. In the first stage he appeals to Nelson Goodman’s ideas about intelligibility (see Goodman, 1954 31-3). Goodman argues that a sceptic about the intelligibility of some purported concept need not provide additional evidence for their view – it is legitimate to follow one’s own ‘philosophic conscience’. The first stage is thus to report the findings of one’s philosophic conscience – that (for Daly at least) grounding talk is unintelligible. The second (and, I think, far more interesting) stage is to rebut the charge of the opposition who say talk of grounding is intelligible.

Both Schaffer (2009) and Rosen (2010) specify the logical properties of grounding, cite connections between grounding and other terms, and give examples of grounding, all by way of trying to make a notion of grounding intelligible. As Rosen (2010: 113) argues, the mere fact that we cannot explain some idiom in more basic terms is not reason to immediately discount it as unintelligible. Many of our ‘best words’ (Rosen’s example here is metaphysical necessity) do not admit of definition, but that does not mean we do not understand those notions. Rosen’s strategy is to engage in grounding talk, cite logical properties and draw out connections, and to see if anything is gained. Daly resists this strategy before it gets off the ground, because he rejects the idea that this work helps to make the notion of grounding intelligible, let alone allows us to assess whether it is useful.

Daly argues that the logical properties of grounding (discussed above) do not fix the content of the term. Like ‘grounds’, ‘explains’ (in the appropriate sense) is (plausibly) transitive, irreflexive, asymmetric, partial, hyperintensional, and non-monotonic. Grounding and explanation are not, however, co-extensional. Causes explain their effects, but don’t ground them (Daly 2012: 91). If notions of grounding and explanation have different contents but the same logical properties, then the logical properties of grounding don’t fix its content. This first of Daly’s arguments sounds like something a proponent of grounding might well accept (indeed, Audi (2012a: 118) does accept it in the course of responding to Daly). By itself, this is no reason to be sceptical about the notion of grounding. Citing the logical and structural properties of grounding restricts the notion so we can get a fairly good idea of what is at stake, even if it does not distinguish grounding from all other notions in the vicinity.

Daly’s second argument (2012a: 91) targets the strategy employed by (amongst others) Schaffer (2009) and Rosen (2010), where analytic connections are traced between grounding and other terms. Daly argues that those other notions are either too close to grounding not to be themselves tainted by its obscurity, or far enough away that their connection to grounding is questionable. Most of Daly’s examples of the former kind concern connections with notions of fundamentality. The latter include connections with notions like explanation (discussed in Chapter 4).

This argument does not function as a knock-down argument against grounding. If successful, though, it robs the grounding-advocate of an attractive way to elucidate grounding talk. The friend of grounding

can still endorse Rosen's plea that we 'relax our antiseptic [sic] scruples for a moment and admit the idioms of metaphysical dependence into our official lexicon', in the understanding that 'if this only muddies the waters, nothing is lost; we can always retrench', but that 'if something is gained...we may find ourselves in a position to make some progress' (Rosen, 2010: 110). We can understand this proposal as distinct from the proposal that we can make grounding intelligible by connecting it to notions of fundamentality. The thought is then that, whether or not we are able to make proper sense of the grounding relation, if we can put the notion of grounding to good use then that is a reason to endorse it. In this form, the argument is a fairly weak one (it seems unlikely that we will be able to find many uses for grounding if we don't really understand the notion), but it might have some success in conjunction with other arguments in favour of grounding, or when fleshed out in greater detail. (I make some contribution to this sort of project in Chapter 5).

Most parties to the grounding debate agree that the most effective way to argue for the intelligibility of grounding is by appealing to purported examples of grounding. Daly (2012a: 95) claims that even this strategy is inconclusive. Anyone who fails to understand 'grounding' will consequently fail to understand any examples using that notion. The sceptic can either claim not to understand the example at all, or can claim of any putative example of grounding that it is really an instance of some other notion. I discuss the latter strategy below. The former seems to introduce a kind of dialectical stalemate. It is true that the sceptic can always deny understanding, and such a denial might sometimes be appropriate. The worry is that such denial is always possible, whether appropriate or not, and if the majority of people think they do have a good enough grip on the notion, the fault may be with the sceptic rather than with the proponent of grounding. If we have a notion that enough people understand enough for it to do useful, recognisable metaphysical work, we at least ought not to dismiss it out of hand.

One further argument in favour of the intelligibility of grounding claims is that they seem not to be exclusively a part of technical, philosophical discourse. Grounding locutions are often offered as part of ordinary language. Thus, we are familiar with the idea that one has a duty to feed one's children in virtue of being their parent, or a right to payment in virtue of having signed an appropriate contract and fulfilled one's obligations. Clark and Liggins (2012: 815) point out that 'makes' and 'thereby' are also

parts of non-philosophical English that at the least are closely related to ‘grounds’ and ‘is grounded in’. Even if it can be maintained that the latter expressions are not part of ordinary discourse, the idea of one thing depending on another in a non-causal way clearly has a role to play in our ordinary lives, and this presents a challenge to the claim that grounding is obscure or unintelligible.

2.2.2. Usefulness

Jessica Wilson argues that there is no distinctive relation of (what she calls) big-G ‘Grounding’. Her view is that there is a plethora of small-g ‘grounding’ relations – relations of type and token identity, part-whole relations, functional realisation, etc., but that these are not (as suggested in Bennett, 2011) determinates of a common determinable. In fact, Wilson argues that even if the more specific relations could be unified under a distinctive type, this would not be reason to posit an additional Grounding relation.

Wilson’s first argument is that Grounding (like supervenience) is too coarse-grained to characterise metaphysical dependence on its own (forthcoming: 5). The case of robustly emergent mental states provides an example. Such mental states are purportedly both fundamental and grounded in the physical. Grounding, Wilson argues, conflates fundamentality and the absence of dependence, and thus cannot give a correct analysis of the phenomenon of emergence (Wilson, forthcoming: 7). (This argument is also made by Barnes, 2012). There are (at least) three ways one might respond to this argument. The first is to simply deny that robust emergence is metaphysically possible, and to tell some other story about what is going on in purported cases of emergence.

The second way to respond to Wilson is to take emergence to be an argument for the priority monist view championed by Schaffer (e.g. 2010a; 2013), rather than an argument against understanding fundamentality in terms of ungroundedness (these issues are discussed further in section 3 below). Schaffer holds that the only fundamental entity is the entire cosmos, and that everything is ultimately grounded in this single entity (see Chapter 3). The phenomenon of emergence (e.g. in entangled quantum systems) provides him with one of his main arguments, because emergent properties are properties of a whole that (arguably) cannot be reduced to properties of the parts (and the relations

between them). Emergence is consistent with a view of grounding that takes the whole to be fundamental.

A third way to respond to Wilson's first argument is inspired by Barnes (2012). Barnes embraces the need for a distinction between the fundamental and the derivative on the one hand, and a further distinction between the ontologically dependent and independent on the other. Barnes takes the notion of fundamentality she works with to be primitive, but attempts to key us into it with examples – the fundamental entities are all and only those God needed to create in order to make the world how it is. Derivative entities are those which are not fundamental – they are always derivative of fundamental entities (not of other derivative entities). Derivative entities are (metaphysically) explained by fundamental entities. Fundamental entities are not explained by anything (Barnes, 2012: 876-7). Rather than thinking it is problematic for grounding that ontological independence and fundamentality can come apart, Barnes claims it as a victory for her understanding of metaphysical structure. She claims that, as a bonus, her account yields a 'natural and attractive' understanding of emergence (Barnes, 2012: 900). It is not clear why proponents of grounding must be wedded to the idea that fundamentality and ontological independence (as characterised by ungroundedness) cannot come apart. I discuss different conceptions of fundamentality in Section 3 below.

Wilson's next argument against 'big-G' Grounding is that philosophers almost never make general grounding claims without a more specific relation in mind (Wilson, forthcoming: 9). For example, when naturalists say that the mental is grounded in the physical, they might be a type-identity theorist, or a token-identity theorist, or a functionalist. When people say that the dispositions of a thing are grounded in its categorical features, they again have in mind either a token-identity theory, or a functionalist theory, and so on. Wilson claims that grounding is 'metaphysically underdetermined' (forthcoming: 9), because further more highly specified accounts of the dependence in question are always available. She argues that it cannot then be the case that Grounding is needed in specific investigations into metaphysical dependence (because we can always work with the more specific account we have in mind).

The best response to this criticism is one Wilson herself considers – that (big-G) Grounding marks an appropriate level of grain for investigations into metaphysical dependence. Grounding is a useful

addition to our toolkit alongside the more specific grounding relations we already admit because it allows us to make appropriately general claims (e.g. that Grounded entities cannot come apart modally from their Grounding entities) (Wilson, forthcoming: 14). We might add that this is cause to reject Wilson's characterisation of things – it is not the case that philosophers always have a more specific relation in mind when they make grounding claims, because sometimes those claims are claims *about grounding*, where grounding is to be understood in Wilson's 'Big-G' sense.

Wilson's reply to this strategy is that it only motivates adopting Grounding as a merely pragmatic, and not as a metaphysical notion. In fact, I am extremely sympathetic to the idea that grounding is a pragmatic or epistemic phenomenon, though I postpone discussion of this until Chapters 4 and 5. Nevertheless, I don't think Wilson has done enough to establish her claim in Wilson (forthcoming). Proponents of grounding can claim that it *is* pragmatic to speak in terms of grounding, but the benefits are not *merely* pragmatic. The friend of grounding can argue that what accounts for the pragmatic benefits of talking in terms of ('big-G') Grounding is that such talk reflects the structure of reality; the notion of grounding is one that 'carves nature at a metaphysical joint'.

A different attack on the usefulness of grounding talk is found in Hofweber (2009). Hofweber defines 'esoteric' metaphysics as metaphysics that is focused on questions involving distinctly metaphysical terms (2009: 267). Idioms of dependence meant in a metaphysical sense are a part of the esoteric metaphysics that Hofweber finds absurd, accusing proponents of such talk of conflating understanding of priority in the sense in which it is familiar from natural language and from more 'egalitarian' metaphysics (i.e. metaphysics where questions are expressed in ordinary, everyday, accessible terms) with a distinctively metaphysical conception of priority. Examples given to elucidate the notion of grounding are of the former understanding of priority, where the notion they are employed to encourage understanding of is of the latter. Purported instances of grounding are really just examples of logical entailment, or conceptual priority, or mathematical priority (see Hofweber 2009: 269).

It is not clear to me that Hofweber is right to denounce all questions couched in distinctively metaphysical terms. Whilst his concerns can serve as a warning against populating one's theory with too many primitives, I don't think it obvious that there are no metaphysical primitives at all. Hofweber says

‘esoteric metaphysics appeals to those...who deep down hold that philosophy is the queen of the sciences after all’ (2009: 273). In fact, all that would be required for there to be questions stated in irreducibly metaphysical terms is that there be some metaphysical questions that do not reduce to questions in another discipline. If in fact there are no such questions, even then we do not show questions about grounding to be ‘absurd’; we just have to do some work to explain in (e.g.) scientific terms what it is we mean by ‘grounding’, and give up on the idea that the notion is a metaphysical primitive. It is not clear that by doing so we destroy the possibility of a notion of grounding doing any interesting work.

Whether or not we buy into Hofweber’s characterisation of esoteric metaphysics, we can resist at least some of his specific attacks on grounding. The idea that we have no use for the term is interesting, but can be responded to along the same lines as we might respond to Wilson. Nothing convincing has been said against the idea that specific examples of particular relations might all fall under a more general notion of grounding, and that that notion might do useful explanatory work, or even carve nature at a joint.

What I find most interesting in Hofweber’s position is the suggestion of a conflation between our ordinary understanding of priority and a philosophical use of the term. In Chapters 4 and 5 I make a similar argument, but I claim that the conflation is between understandings of the notion of explanation rather than priority.

2.2.3. Indispensability

In the course of responding to sceptics about grounding, we have seen that grounding is often defended or motivated by the contention that we can give plausible examples of grounding (e.g. Correia and Schneider 2012; Schaffer 2009; Rosen 2010); cite connections between grounding and other notions (e.g. Schaffer 2009; Trogdon 2013b); and use grounding to do useful philosophical work (e.g. Fine, 2001; Rosen, 2010; Schaffer, 2010a). Audi (2012a; 2012b) adds to the arsenal of friends of grounding by offering a more explicit form of an indispensability argument for grounding.

Audi's argument concerns the distinctive form of non-causal explanation that tracks grounding relations (at least on Audi's understanding of the connection between grounding and explanation – see his 2012a: 120, and Chapter 4 of this thesis). His argument is as follows (2012a: 105):

- (1) If one fact explains another, then the one plays some role in determining the other.
- (2) There are explanations in which the explaining fact plays no causal role with respect to the explained fact.
- (3) Therefore, there is a non-causal relation of determination.

The argument is valid, but Audi accepts that the premises are controversial. Concerning the first premise, Audi suggests that one might argue first that it makes too robust a notion out of explanation, and second that determination is not the only relation that might account for a successful explanation (2012a: 105). The former strategy I think represents a serious worry. My concerns about making explanation too robust are explored in Chapter 4 (where I suggest taking grounding to be a form of explanation where explanation is to be understood as having some agent-relative features) but I set them aside for now. Certainly the majority of the literature understands grounding as a robust form of metaphysical explanation, and much needs to be said to argue against this conception.

Motivating his second worry about the first premise, Audi has in mind identity, where, for example, the identity relation that obtains between water and H₂O accounts for the way in which the fact that there is H₂O in the glass explains the fact that there is water in the glass. If this is right, then there can be explanations without a determination relation. Audi's response is to claim that if the relation here is really identity, then it cannot also be explanatory. Audi claims that explanations require both asymmetry and irreflexivity, but identity is both reflexive and symmetric. Audi argues that the relation between the water and the H₂O in the glass is therefore not identity, but constitution, such that the H₂O constitutes the water. Since constitution has the required logical properties, it can be explanatory. Since it is non-causal, constitution is a non-causal relation of determination, and Audi's argument goes through. Note that even if Audi's opponent can deny that constitution is explanatory (and thus deny that it is an example of a non-causal determination relation) not every apparent case of non-causal determination can be

accounted for by appeal to constitution instead (Audi, 2012a: 106). Audi's opponent will have to argue that no seeming example of non-causal explanation is a case involving non-causal determination.

Audi's second premise appeals to examples to motivate the claim that there can be non-causal explanations. The examples he gives are grounding claims – for example, that normative facts are grounded in non-normative facts, and that semantic facts are grounded in social and psychological facts (2012a: 106). The examples themselves are fairly convincing – the non-normative facts don't cause the normative facts, they just explain them – but it is perhaps problematic to use the terminology of grounding here, when this is an argument for the need for a grounding relation. Audi might instead use the locution 'explains' and then argue that we ought to introduce the term 'grounding' in order to distinguish causal from non-causal determination. A more serious concern is that by using grounding terminology in support of the second premise, Audi's argument becomes epiphenomenal. The real work is done by introducing compelling examples of grounding, and this is a strategy we have seen multiple times by way of justifying grounding talk.

2.3. Connections

Whilst grounding is generally taken to be primitive, it is often explored via connections with other notions. Here I briefly discuss connections between grounding and modality, essence, truth-making and reduction. I save connections between grounding and notions of metaphysical structure and fundamentality for the next section. Connections between the notion of grounding and that of explanation are explored in detail in Chapter 4, and so I do not discuss them in this chapter.

2.3.1. Modality

Part of the motivation for accepting a notion of grounding into our 'philosophical toolkit' is that modal notions are too coarse-grained to capture what is at stake in arguments about dependence (see sections 1.2 and 2.1). Whilst grounding is not an entirely modal notion, it does have modal consequences. There is thus a question about precisely how more familiar modal notions interact with our notion of grounding. Correia and Schnieder (2012: 20) suggest that a common idea is that if some fact f is grounded in some facts g, g', \dots , then when g, g', \dots obtain, f obtains as a matter of necessity (when g, g', \dots fully

or completely ground f). If, as Fine suggests (and I mentioned above) we can distinguish between different kinds of grounding corresponding to different species of necessity, then we can understand the above principle as involving conceptual or logical necessity instead of metaphysical necessity.

Trogon (2013a) raises a problem for the distinction between full and partial ground discussed above (section 2.1) which is relevant here. We have said that when g, g' fully ground f , and when g, g' obtain, f obtains as a matter of necessity. But the distinction between full and partial ground does not properly correspond to a modal distinction, because there can be facts corresponding to necessary truths (and thus obtaining as a matter of necessity) that have merely partial grounds. Suppose A and B are both necessary truths. The fact that $A \wedge B$ is only partially grounded in the fact that A, even though the fact that $A \wedge B$ is true in every possible world. Trogon takes this to demonstrate that the distinction between full and partial grounding is an explanatory rather than a purely modal distinction (2013a: 467). This further supports the idea that whilst modality and grounding are connected, the latter cannot be explained in terms of the former.

2.3.2. Essence

Fine (2012b) argues for a connection between essence and ground that has its roots in the idea that when some fact f is grounded in $g, g'...$, this has something to do with the nature of the entities in question. This connection is not, however, just that it lies in the nature of some fact f (or the entities it involves) that it should be grounded in the facts $g, g'...$ given that those facts obtain. This is inferred from cases like that of existential generalisation, where even though the fact that someone is a philosopher is grounded in the fact that Socrates is a philosopher, the fact that someone is a philosopher 'knows nothing of Socrates' (Fine, 2012b: 75). It therefore does not lie in the nature of the fact that someone is a philosopher that it be grounded in Socrates if he exists (or in the fact that Socrates exists if he does).

Fine rescues the connection between essence and ground by generalising the statement of ground. It does sound plausible to say that it lies in the nature of the fact that someone is a philosopher that 'this fact will, for any person x , be grounded in the fact that x is a philosopher given that x is indeed a philosopher' (2012b: 75). The explanatory connection that (for Fine at least) lies at the heart of

grounding claims is itself explained in terms of the nature of the grounded fact, rather in terms of the nature of the grounding fact. In other words, it is of the essence of things that they be grounded as they are. This (as Fine, 2012b: 76 notes) has interesting consequences for the methodology of grounding claims. In answering questions about what grounds what, we should look to the grounded fact (or entity), and not to the grounding facts (or entities).

Rosen (2010: 132) rejects the kind of link between ground and essence proposed by Fine. One of his examples motivating this rejection concerns a non-reductive materialist in the philosophy of mind. Our materialist claims that all of the phenomenal facts are grounded in the physical facts. She states, for example, that pain is grounded in something like c-fibres firing. She would deny, however, that there is anything in the nature of pain that requires that it be grounded in the firing of the c-fibres. More generally she denies that it is in the nature of pain that it be grounded in brain states at all. The nature of pain is entirely exhausted by its phenomenal character, even though it is always the case that I am in pain in virtue of some fact about the state of my brain.

Fine's response (2012a: 77) is to argue that that Rosen's counterexamples conflate different conceptions of ground. The above example is only plausible when we read grounding under a natural rather than under a metaphysical conception. Fine says that the view that the mental is grounded in the physical is only plausible for the natural conception of ground. When the conception of ground we have in play is not metaphysical, we should not expect the grounding facts to hold in virtue of the essentialist nature of the entities involved. This might work for Fine, but many others (including Rosen) take the relation that purportedly obtains between phenomenal facts and physical facts to be a canonical example of metaphysical grounding. Whether we side with Fine or with Rosen, it is understandable that no complete analysis of the notion of ground in terms of essence is available, either because metaphysical grounding is not the whole story, or because there are cases of metaphysical grounding that cannot be analysed in terms of essence. As with modality, it seems that ground and essence are connected but distinct.

2.3.3. Reduction

Rosen (2010: 122) argues for a connection between grounding and reduction such that when some fact f reduces to some fact g , and f is true, then f is grounded in g . By 'reduces to', Rosen has in mind the idea

that one fact consists in another. For example, the fact that $\neg\neg P$ reduces to the fact that P because the fact that $\neg\neg P$ consists in the fact that P (when P is true). If it is the case that to be a vixen is to be a female fox, then Sandy is a vixen in virtue of being a female fox. But this presents us with a puzzle, similar to the issue we ran up against in defending Audi's first premise of his argument for grounding: if we have correctly defined the term 'vixen', then the fact that Sandy is a vixen and the fact that she is a female fox are the very same fact. If the latter fact nevertheless grounds the former, then we have a violation of the principle of irreflexivity. This is also the same problem as we encountered before in discussion of Jenkins' argument against the irreflexivity of grounding, but here the target is the grounding-reduction link rather than the irreflexivity constraint on grounding. Rosen (2010: 124) argues that to retain the link, reduction must be a relation between two distinct propositions – we can't characterise reduction in terms of identity, and so we ought not to express reduction with a 'to be F is to be G ' locution.

This, in fact seems to be how we do think about reduction. It sounds odd to say that Jonas' being an unmarried male consists in (or reduces to) his being a bachelor, but fine to state that his status as a bachelor consists in (reduces to) his being an unmarried male. This asymmetry seems to be an explanatory one – we think that Jonas is a bachelor in virtue of being an unmarried male because we understand the term 'bachelor' to mean 'unmarried male'. We don't (usually) come to understand what an unmarried male is when somebody explains that it is a bachelor.

Rosen (2010: 124) understands reduction as a kind of metaphysical analysis – of some simple proposition into its real definition. His example (2010: 125) concerns the reductive analysis of being a square; what it is to be a square is to be an equilateral rectangle. The proposition <something is a square> reduces to <something is an equilateral rectangle> (see also Trogdon, 2013b: 117). Reducing some fact f to another fact g thus involves reducing a relatively simple, perhaps unstructured proposition to a distinct, relatively complex, structured proposition. We can then claim that the truth of the more complex proposition grounds the truth of the simpler proposition (Rosen, 2010: 26).

2.3.4. Truth-making

We have come to be familiar with the invocation of truth-making to connect a truth-bearer and whatever accounts for the truth (the truth-maker). The truth is thought to be true in virtue of the existence of the truth-maker, and so an equivalence has been proposed between grounding and truth-making: x makes a proposition y true iff x 's existence grounds the truth of y . (This can be modified to accommodate an operational conception of grounding.) So, for example, it is snow's whiteness that grounds the truth of <snow is white>.

Cameron (2010: 184) finds it natural to define ontological dependence in terms of truth-making, such that x is ontologically dependent on the y s iff the y s, and not x itself, are truthmakers for < x exists>.⁴ Liggins (2012) argues against integrating truth-maker theory into a general account of ontological dependence. His discussion is motivated by the idea that the heart of the truth-maker theory concerns truth's dependence on reality – a relation which propositions bear to other things (2012: 266). He considers other general theories of ontological dependence, in particular that he calls 'the fact-fact theory' (which takes all instances of ontological dependence to be a relation between two facts) and argues that they can accommodate the dependence of truth on reality, and we do not have good reasons to rule them out. Theories like the fact-fact theory are incompatible with truth-maker theory (which relates propositions to other things). Whilst the fact-fact theory can be extended to form a general account of ontological dependence, Liggins argues that the truth-maker theory cannot.

An extension of the truth-maker theory would construe all instances of ontological dependence as making relations, so that (for example) being willed by the gods is the right-maker for a given act, or the arsenic is the poison-maker for the cup of tea. Ontological dependence claims should then be delineated by a three-place predicate ' x is made F by ϕ ' (Liggins, 2012: 268). The problem, Liggins argues, is that there are cases of ontological dependence this account cannot capture. Consider, for example, that grass is not-red in virtue of being green, or that necessarily, water contains hydrogen in virtue of the essence of

⁴ I'll not unpack Cameron's proposal here, since it would involve a discussion of his account of truth-making that is beyond the scope of this thesis

water. Liggins takes these (and numerous other examples) to confirm that there is more to ontological dependence than making.

Like Liggins, Fine argues that the truth-maker theory has ‘an unduly constricted conception of what is grounded’ (2012a: 43). He grants that by looking for truth-makers we may discover genuine grounds, but that the truth-making relation is merely ‘a pale and distorted shadow of the notion of genuine interested to us’. Truth-maker theory is only concerned with the grounds for the truth or correctness of our representations of the world. Whenever we consider the question of what makes some representation that P true, there is a further question about what makes it the case that P, which truth-maker theory is not equipped to answer (Fine, 2012a: 43). Fine is also concerned (like Liggins) by the way in which truth-making dictates that the ultimate source of what is true always lies in what exists. There seems to be no reason that truth could not depend on something relational (i.e. on x standing in relation R to y , or the negation of something relational, or something of another form (Fine, 2012a: 44).

A different concern for truth-making addressed by Fine concerns the truth-making relation itself. Truth-making is usually expressed in modal terms (so f is a truth-maker for P iff the existence of f necessitates the truth of P) but as we have seen, modal theses are too coarse-grained to express grounding claims. Anything can be a truth-maker for any necessary truth, and apparent grounding asymmetries between necessary co-existants (like Socrates and {Socrates}) are not captured by the view. Just as the fact that Socrates exists is truth-maker for the proposition that {Socrates} exists, the fact that {Socrates} exists will be a truth-maker for the proposition that Socrates exists (Fine, 2012a: 45).⁵ We therefore have a number of good reasons to think that the connection between ground and truth-making is not one of identity, but we might instead wish to account for what is at stake in truth-making claims in terms of a more general notion of grounding. If truth-making is to be accounted for in terms of grounding, then truths are grounded in their truth-makers. Thus, the truth of the proposition <snow is white> is grounded in something like the state of affairs of snow’s being white, or in the fact that snow is white.

⁵ One way to respond to this would be to propose an account of making that builds in some kind of relevance constraint, though it is not obvious how this would be best articulated.

3. Fundamentality

We have thus far skirted around the connection between ontological dependence and grounding, and notions of fundamentality. The subject of this thesis is metaphysical structure, and how we might understand that structure (if indeed reality has a privileged structure)⁶. One aspect of this concerns which things are fundamental; what are the metaphysical foundations of reality. There are multiple ways to cash out ideas about fundamentality, almost all of them intimately related to notions of dependence. In most cases, the notions that are used to characterise fundamentality can be employed to characterise a richer metaphysical structure. In this section I attempt to separate these different conceptions of fundamentality (though note that in much of the literature two or more of them are run together.)

I use x to stand for any type of entity we might take to be fundamental, be it fact, property, concrete object, set, or otherwise. Proponents of each of the approaches to fundamentality I discuss might take the bearers of fundamentality to be a limited class of those things. Since this has little bearing on the attempt I make to distinguish conceptions of fundamentality here, for the most part I ignore those nuances.

3.1. Ungroundedness

Perhaps the most common way to understand fundamentality is via grounding (or more generally, ontological dependence). Schaffer (2009: 373) offers the following principle:

- (a) x is fundamental \equiv_{df} nothing grounds x

The intuitive idea is that anything that itself depends on something else cannot be fundamental. The picture this generates is one where chains of grounding (adhering to the structuring principles of asymmetry, irreflexivity, and transitivity) terminate in an entity or collection of entities that has no grounds. This account of fundamentality takes an absolute notion of fundamentality (i.e. x is fundamental) rather than a primitive notion of fundamentality to be basic.

⁶ Eli Hirsch (e.g. 1993) is a prominent defender of the broadly Carnapian view that reality does not have a privileged structure.

A defender of this conception of fundamentality might be happy to stop there, but it is very common to hold that fundamentality comes in degrees. Ascriptions of greater and lesser degrees of fundamentality might then be given relative to this absolute notion, though it is a little tricky to give an adequate account of how this should work. At first blush we might be tempted by a proposal such as the following:

(b) x is more fundamental than $y =_{df}$ x grounds y

Recall that grounding is usually taken to be a strict partial ordering on entities. It is tempting to think that if x is more fundamental than y then x must ground y , but the atoms that compose my table are plausibly more fundamental than my computer, even though the table atoms don't ground the computer. We might try to fix the problem by emphasising that atoms comprise a 'level of reality' that is closer to the fundamental than complex objects like computers and tables, but it is not hard to think of examples where separate grounding chains appear incommensurable. The question of whether the set whose sole member is Wilfred the 18th Century water molecule is more or less fundamental than my table admits of no easy answer. Grounding is only a partial ordering – not every pair of entities in the relevant domain is related by the grounding relation.

Perhaps we could simply write out the chains of grounding connecting the derivative entity in question to some fundamental entity, and compare the lengths of the chains in order to establish claims of relative fundamentality (this is similar to the Lewisian method for assessing the naturalness of a given property – see his 1983; 1984, and Chapter 2 of this thesis). But go back to our example. The table is grounded in the atoms that compose it, but is it first grounded in some bigger parts of itself (e.g. its legs and its top)? {Wilfred} is grounded in Wilfred, but what grounds Wilfred? Perhaps the atoms that once composed him? But those are presumably now scattered across the globe (maybe one or two of them is even in my table). Perhaps Wilfred is instead grounded in whatever grounds facts about the past (the present perhaps), and in turn in whatever grounds that. Either way, it seems as though incommensurability remains a problem in giving accounts of relative fundamentality between entities that have little to do with one another. Perhaps the best we can do for now then is to say that x is more fundamental than y if x grounds y . This is a far weaker claim than the definition in (b), and seems to fall short of an adequate account of relative fundamentality because it tells us nothing about when the condition does not obtain.

A problem with understanding fundamentality in the way described in (a) is that we might not want our account of fundamentality to derive from that of grounding. As we saw in section 2.2, there is a core of resistance towards adopting a notion of grounding as part of our metaphysical practice. It might be that notions of fundamentality are important and useful, but that the grounding debate is eventually won by the dissenters. If that were the case, and if fundamentality were to be explicated in terms of grounding, then we would be in danger of throwing the baby out with the bathwater. For this reason, we might wish to resort to taking whatever is fundamental to be that which is non-dependent (understood in some way that does not appeal to grounding) rather than ungrounded.

A further worry about taking what is to be fundamental to be that which is either ungrounded or non-dependent is that it immediately rules out the possibility of fundamental dependent entities. This consequence might be more troubling than we first suppose. Firstly, it precludes any fundamental entity from being self-dependent. Defenders of arguments for the existence of God such as the cosmological argument often claim that God is the sufficient reason for God's own existence, or that God's nature demands God's existence. If God is a fundamental entity (as is supposed by defenders of those arguments) then either the present characterisation of fundamentality must be rejected, or an alternative explanation for why we need not look for a cause for God must be provided.

Schaffer argues that fundamental entities must be modally unconstrained, and open to recombination. Since the fundamental elements of the world serve as building blocks for reality they must be metaphysically independent, so that 'any combination of ways that each entity can be individually is a way that the plurality can be collectively' (2010a: 40). This, together with the requirement that the basic entities collectively cover the cosmos (providing a 'blueprint for reality') generates a constraint on possible answers to the question of what the basic actual concrete objects are. This constraint Schaffer calls 'the tiling constraint' – basic actual concrete objects must collectively cover the cosmos without

overlapping.⁷ This constraint supports a conception of the fundamental as the ontologically independent, but there might be reasons to resist the tiling constraint.

As mentioned in section 2.2.2, Barnes (2012) argues that the best way to characterise the phenomenon of emergence is to take the distinctions between fundamentality and ontological independence to cut across each other – ontological emergence is the thesis that there are some fundamental things that are not ontologically independent. Emergent entities are not things that ‘come for free’, but must be mentioned in a complete description of the world. This suggests they are fundamental. Nevertheless, emergent entities are sustained by other entities, in the sense that they require, at each moment of their existence, the existence of something else. (Consider, for example, the way in which mental properties might be said to emerge from physical properties. We don’t get the mental properties for free once the physical properties have been specified, but without the physical properties we don’t get the mental properties at all.) Whilst there are alternative ways to characterise emergence, Barnes claims that her suggestion is compelling in its simplicity, and avoids some standard problems associated with emergence (see Barnes, 2012: 892-9). The view that there is any such phenomenon as emergence in the world is a subject of fierce debate, though we should note that the metaphysical possibility of emergence is reason enough to be wary of an analysis of fundamentality that cannot account for it.

One more reason to separate ontological independence and fundamentality is to maintain the possibility of interdependent fundamental entities. In Chapter 3 of this thesis, I argue for a conception of metaphysical structure whereby the asymmetry constraint on grounding is dropped. Though in the development of that position I do not take any part of reality to be fundamental, a related position allows interdependence at the fundamental level. Imagine, for example, a descending hierarchy of levels of reality characterised by derivative entities (the *y*s) related in chains of grounding, terminating in some basic level of entities (the *x*s), where all of the *y*s depend on the *x*s, and all of the *x*s depend on the other *x*s (but not on any of the *y*s). This sort of picture might seem plausible to those who take seriously the

⁷ Though Schaffer’s argument concerns actual concrete objects, there is an obvious generalisation of the tiling constraint to any entities that are candidates for being the bearers of fundamentality. For example, Lewis (1986b: 60) claims of his perfectly natural properties that there must be ‘only just enough of them to characterize things completely and without redundancy’. This is discussed further in Chapter 2, section 2.2.

idea that at the ‘bottom level’ of reality are quantum entangled particles or quantum field values. A definition of fundamentality characterised by independence rules out anything like this view.

3.2. Ontological commitment

A second way to understand fundamentality has its roots in Armstrong’s claim that supervening entities are an ‘ontological free lunch’. The idea is that derivative entities are ‘no addition to being’ (Armstrong, 1997: 12) – an ontology is no less sparse for containing derivative entities than it is for containing the entities that ground them. Rather than measuring our ontology against the Quinean criterion of what a theory is committed to (i.e. against whatever is the value of a bound variable in our best theory of the world), we should instead be concerned with what is fundamental; rather than sticking to the Ockhamist maxim of not multiplying entities beyond necessity, we ought instead to be committed to not multiplying *fundamental* entities beyond necessity. We can thus define a fundamental entity in the following way:

(c) x is fundamental \equiv_{df} x is ontologically committing

Or, perhaps more naturally, with the notion of a derivative entity as basic:

(d) y is derivative \equiv_{df} y is not ontologically committing

The main problem with taking ontological commitment to be the mark of fundamentality is that it leaves open what it is for an entity to be something to which we are ontologically committed. ‘Ontological commitment’ is a technical term. Answers to the question of what it is to be ontologically committed to something are likely to serve just as well as answers to the question about how to distinguish fundamental entities. Though the view that derivative entities are ontologically uncommitting is fairly widespread, this is not often given as a way to mark the distinction between the fundamental and the derivative, rather it is a consequence of that distinction. Amongst those who welcome that consequence are Schaffer (2009) and Sider (2011). Fine (2001) and Cameron (2008) take the distinction between the fundamental and the derivative to be a distinction between what really exists and what ‘merely’ exists. Cameron then explicitly equates his ontological commitment with what really exists, and Fine seems to be motivated by a similar thought. Built into this conception is the idea that fundamental

reality should be complete, in that reality is completely specified once the fundamental facts are fixed. Recall Schaffer's phrase – the fundamental entities provide a 'blueprint for reality' (2010: 39).

An analysis of fundamentality that is in the same spirit as that suggested here is to take to be fundamental whatever entities God would need to think in terms of when creating the world. Once those elements are fixed, the rest follow. This sort of idea is considered in Sider (2011: 138). Barnes (2012: 876) suggests that the creation metaphor is a helpful way to get a handle on fundamentality talk, but will not serve as an analysis of fundamentality. The reason is that we ought not to rule out by definition the idea that the fundamental elements of the world might be necessary, or that there might be necessary connections between fundamental existents. This principle is in direct conflict with Schaffer's thought that the fundamental elements ought to be open to free modal recombination.

3.3. Primitivness

The approach Barnes favours is to take fundamentality as primitive, and hence not to offer any kind of analysis of the concept. Instead, she appeals to metaphors such as the creation metaphor discussed above to elucidate the idea. Another way to get a handle on primitive fundamentality is to take fundamental entities to truth-make their own existence, and to serve as truth-makers for other entities (again, this gloss is similar to the idea that derivative entities are no addition to being over the fundamental).

Taking fundamentality to be primitive incurs a theoretical cost – it is something that we accept that our theory cannot explain. In general, the fewer primitives we adopt the better, because the less we then leave unexplained. Whilst any theory must presumably take some things as primitive, we need good reason to think that fundamentality is among them. Given the numerous other ways to understand fundamentality discussed here, we need arguments for why each of the alternatives is lacking.

A further worry is that taking something as primitive leaves it open to sceptical attack. As in the case of grounding (see Daly, 2012), there is a worry that if we cannot give a clear definition or an analysis of fundamentality, there is no reason to believe that any one notion is picked out by the concept. The term has not been properly introduced, and as such its intelligibility can be called into question. The way to

respond to this is to gesture towards what is meant by fundamentality (as Barnes does) by the use of examples, but it is always open to the sceptic not to understand the notion they are intended to shed light on. (As Lewis (1986b: 203, note 5) quips, ‘any competent philosopher who does not understand something will take care not to understand anything else whereby it might be explained’).

3.4. Linguistic conception

Williams (2012) argues that fundamentality is not itself a metaphysical notion, but rather that it is a device used to express a metaphysical notion. Like those who take ontological commitment to mark the distinction between fundamental and derivative, Williams is motivated by the desire to reconcile a minimal ontology with an endorsement of common sense belief. Williams takes fundamentality to be a linguistic device that enables us to talk about the contents of reality. Fundamentality is an operator that attaches to propositions, where the idea (very broadly) is that we cannot reliably read the ontological commitments of the world from their apparent subject matter. Mind-independent reality is represented by us in thought and language, and it will (presumably) ultimately turn out that some of the propositions expressed by those representations will be true. However, the claim ‘there are tables’ might be held to be true without requiring of reality that there be tables. Instead, what we require might merely be commitment to the truth of the proposition <there are simples arranged table-wise>. In order to mark the distinction, we make use of the operator ‘fundamentally’. This conception of fundamentality can be roughly stated as follows:

- (e) Fundamentally, $P =_{df}$ the truth of P directly requires that reality be a certain way.

Propositions that are fundamentally the case support other true propositions – the sentences that express the propositions that are fundamentally the case have a surface form that directly reflects the requirements on reality for the truth of that proposition. Not always so for the propositions that are not fundamentally the case. Williams describes his proposal as describing a function f from sentences S of natural language into some metaphysically revealing language. The account then says that: for all S , what is required of reality for S to be true is for $f(S)$ to hold (2012: 176). The term ‘fundamentally’ is a signal that we intend to speak in the metaphysically revealing language.

Williams' account includes a framework for specifying the relevant requirements on reality. Williams takes these reality requirements to be compositional (2012: 180). This is in contrast to the paraphrase strategy favoured by Quine, where specifications are given for whole sentences rather than individual components. Williams favours a metasemantic approach of the kind developed by Lewis (e.g. 1984) and discussed in the next chapter, and in more detail in the appendix. The core of this approach is that it maximises both simplicity on the one hand, and fit with a certain set of privileged correlations between sentences and propositions on the other (Williams, 2012: 181). This trade off allows Williams to give an account of requirements on reality for the truth of propositions whereby those requirements are weaker than they appear to be given the surface grammar of the sentence expressing the proposition. If the statement 'there are tables' required that there really be tables, then in a nihilistic world (where all that really exists are simple particles arranged in various ways) the vast majority of people would be speaking in error. It is thus more charitable to interpret that statement as requiring instead that there be simples arranged table-wise. Fundamentally, there are simples arranged table-wise.

This conceptual gloss on fundamentality is markedly different from the purely metaphysical accounts discussed above. Note also that it is incompatible with a description of entities such as objects or events as fundamental, at least without some modification. (Though one such modification would take the claim 'x is fundamental' to convey something like 'fundamentally, there are x's', where Williams' treatment is then given to the latter sentence.)

3.5. Naturalness/structure

Lewis (1983) introduces the idea that there is an elite class of perfectly natural properties and relations that 'carve nature at the joints'. One way or another (see the next chapter for details) naturalness admits of degree, and properties can be ranked in accordance with their relative degrees of naturalness. At the limit are the perfectly natural properties, and Lewis implies that these properties are fundamental. This idea is picked up in much of the literature, where 'perfectly natural' and 'fundamental' are taken to be synonymous (cf. Dorr and Hawthorne, 2013: 37). This suggests a fifth characterisation of fundamentality:

(f) x is fundamental \equiv_{df} x is perfectly natural.

In the next chapter I discuss naturalness theory in detail, and so I will not give an account of it here. For reasons discussed in the next chapter to do with the ability of the property of perfect naturalness to underwrite the role it is employed to play, I do not think that there is a viable notion of perfect naturalness, and so I think this characterisation of fundamentality fails.

Sider (2011) expands and extends the notion of naturalness into a wider theory of what he calls ‘structure’. For Sider, the heart of metaphysics concerns what reality is fundamentally like, and fundamentality is a matter of structure – of carving the world at its joints (Sider, 2011: 5). His conception of fundamentality is thus:

(g) x is fundamental \equiv_{df} x is part of the world’s structure.

Sider’s notion of structure is a metaphysical primitive, and as such it resists definition. It is closely related to the idea of a fundamental language – a language whose predicates, but also quantificational structure and operators carve nature at the joints. Sider (2011: 92) posits a primitive ‘structural’ operator that attaches to any portion of a language, and maps portions of language to truth iff that portion of language fits the fundamental structure of reality. Because this idea of a fundamental structure or privileged carving of reality is central to Sider’s proposal, his views about fundamentality are, (unlike Williams’), very much metaphysical.

Some of the criticisms of naturalness theory that I make in the next chapter also apply to Sider’s view (see Chapter 2, section 5). A different worry is Sider’s admission that a ‘knee-jerk realism’ about structure is a presupposition of his position. Sider counts this unflinching conviction in a realism he declines really to argue for his ‘deepest philosophical conviction’ and confesses that he has ‘no idea how [to] convince someone who didn’t share it’ (2011: 18). For my part, I find myself as baffled by his conviction as he is by those who lack it. But since an exchange of baffled looks is no substitute for philosophical argument, I do not pursue this line. A more concrete problem for Sider is addressed by Schaffer (2014), and concerns a mismatch between the roles Sider identifies that motivate the need for concepts of ‘more structural than’ and ‘structural enough’, and his primitive notion of absolute structure. I discuss this problem in Chapter 2.

3.6. Truth-makers

The penultimate conception of fundamentality I'll distinguish here is one we have already touched on as a gloss for understanding primitive fundamentality, but I wish to discuss it in a little more detail. We can define it as follows:

(h) x is fundamental \equiv_{df} x is a truth-maker

As presented in Cameron (e.g. 2010: 184) the idea is that the truth-makers make true all of the true sentences of our best theory. These truth-makers are therefore the fundamental entities. (This is closely related to conceptions (c) and (d) above.)

We saw in section 2.3.4 above that there are good reasons to think that truth-making and grounding are distinct. The criticisms that Liggins and Fine employ against truth-making as a theory of ontological dependence will also apply when considering truth-making as a theory of fundamentality. Truth-makers must be objects or existential facts, and so a truth-making conception of fundamentality promises a (qualitatively) sparse fundamental landscape. Though this might be seen to be an attractive feature of the view, it rules out (perhaps too quickly) the possibility that non-objects or non-existential facts might be amongst the fundamental things.

A further problem is discussed by Sider (2011: 160), who argues that truth-maker theory is explanatorily unsatisfying. If what is fundamental are truth-makers (i.e. either entities, or facts of the form ' x exists') then ultimate explanations will always terminate in citations of entities. Such lists are so unstructured that they cannot provide satisfying explanations. Sider claims that 'there's something wrong with a fundamental theory of reality that precludes satisfying ultimate explanation' (2011: 160). He argues that unlike other views of fundamentality, the truth-maker view does not give an explanation of the relation between the fundamental and the non-fundamental beyond stating that the former makes-true truths about the latter. There is no scope for degrees of fundamentality on the truth-maker view.

A response to this particular criticism might be available to the truth-maker theorist. As with definition (a) above, we can imagine an extension of (h) to relative fundamentality. We might, for example, take the entailments E of sentences that some truth-maker x makes true to be less fundamental than the

sentences S that x itself makes true. Again, however, ascriptions of relative fundamentality based on entailment from different truth-makers seem incommensurable. An account of relative fundamentality on the truth-maker view might be available, but not entirely satisfying. Proponents of the truth-maker view of fundamentality will also inherit the traditional problems of truth-maker maximalism (the view that every truth is made-true by some truth-maker), such as accounting for negative existential truths for which there is no obvious truth-maker (see MacBride, 2014).

3.7. Explanation

Sider's above criticism of the truth-maker view raises a point that I think ought to be central to our understanding of fundamentality – fundamental truths ought to be explanatory. A final approach to fundamentality makes the connection between explanation and fundamentality explicit:

- (i) x is fundamental =_{df} x provides an ultimate explanation.

In Chapter 4 of this thesis I discuss explanation in detail, but here I leave it deliberately open what kind of explanation is required. Debates about fundamentality are generally concerned with metaphysical explanation (the topic of Chapter 4), and so whatever is fundamental can be considered that which provides the ultimate (metaphysical) explanation. This is an idea that seems to permeate the work of Schaffer (e.g. 2009, 2010a), Sider (2011) and Fine (e.g. 2001, 2012a) amongst others. I think there might be other interesting notions of fundamentality understood (at least partly) in terms of explanation and couched in less metaphysical terms, but again, I postpone discussion until later.

Given that our judgements about explanation tend to be relative rather than absolute, an intuitively attractive option here would be to take a relative notion of fundamentality to be basic, and to define a notion of absolute fundamentality out of the relative one. Perhaps the most obvious way to define a notion of relative fundamentality is as follows (for some suitable notion of explanation):

- (j) x is more fundamental than y =_{df} x explains y .

Unfortunately, this account inherits the sort of problems that arose for the analogous definition of relative fundamentality in terms of grounding; an electron in Germany is presumably more fundamental than my finger, but doesn't explain it. It's not clear how the account can be modified to avoid this

problem. In general, taking relative notions of fundamentality to be basic is a possibility that is rarely discussed in the literature except in terms of naturalness. I discuss the possibilities for formulating naturalness theory in Chapter 2.

4. Concluding Remarks

In this chapter I have introduced the notions of ontological dependence, grounding, and fundamentality with which this thesis is concerned. I have provided a survey of the structural features of grounding and connections between grounding and other notions, and discussed some reasons both for and against adopting a notion of grounding in metaphysics. I have attempted to distinguish a number of different conceptions of fundamentality, and highlighted some of their strengths and weaknesses.

Of these conceptions of fundamentality, by far the most popular (at least in the recent literature) are conceptions that appeal to naturalness or to grounding. Both of these conceptions of fundamentality are thought to have the resources for characterising relative as well as absolute fundamentality, and so are *prima facie* suited to limning the structure of reality. The rest of my thesis will therefore be devoted to a critical examination of those two notions. In the next chapter I focus on naturalness, and also make some remarks about Sider's own theoretical posit, structure, which he takes to be an extension of naturalness. The focus of the final three chapters will be grounding. In particular, I am interested in the relative notion of fundamentality; in what it is for one thing to be more fundamental than another. In some cases (as with naturalness, as I explain in the next chapter) it is possible to define a relative notion of fundamentality out of an absolute notion. In others, as with grounding, this is problematic (see section 3.1), and we can focus instead on what it is for one thing to ground another and use this to characterise reality's structure.

I will have little more to say about the truth-making conception of fundamentality, or much that is explicitly about either the primitive or the ontological commitment conceptions, though the some of the motivations for adopting these accounts of fundamentality will be discussed in the context of alternative accounts.

Chapter 2: Naturalness and Structure

The focus of this chapter is the notion employed in the fifth characterisation of fundamentality introduced in the previous chapter: x is fundamental \equiv_{df} x is perfectly natural. One of the best candidates for understanding the structure of reality is to identify the fundamental with the perfectly natural properties; perfectly natural properties are thought to ‘carve nature at its joints’. Ascriptions of relative naturalness can be taken to characterise the rest of reality’s structure.

In section 1 I explain the notion of naturalness. In section 2 I go through some of the key roles that the perfectly natural properties have been thought to play, and argue that *perfect naturalness* itself largely fails to play the relevant roles. Section 3 assesses the relative degree of naturalness had by the property *perfect naturalness*, in the light of the conclusion of the previous section – that *perfect naturalness* is not perfectly natural. Section 4 is a discussion of the consequences of this argument, and in section 5 I shift focus to Sider’s extension of naturalness – structure, and argue that that notion is affected by a related problem. Neither naturalness nor Sider’s structure are suited to characterising the structure of reality.

1. Naturalness

David Lewis’ (1983) ‘New Work for a Theory of Universals’ motivates and develops a distinction between ‘perfectly natural’ properties and relations,⁸ and the rest. (Perfectly) natural properties are held to be simple, intrinsic, and non-gerrymandered. The distinction between these and the less-than-perfectly-natural properties is put to work in many different ways, including helping to solve problems with induction and reference (see appendix), and giving accounts of phenomena as diverse as duplication

⁸ Henceforth in this chapter the term ‘property’ should be taken to include relations.

and lawhood. The many proponents of naturalness theory take the natural properties to play a key role in metaphysics.

Naturalness theory is posited as a unifying solution to a variety of metaphysical puzzles. The inclusion of a notion of naturalness in the ideology of its defenders is justified by appeal to its theoretical utility and import (compare the arguments made by e.g. Schaffer (2009) and Rosen (2010) and discussed in the previous chapter; that a notion of grounding is required because of the various theoretical roles the notion can play). Defenders of naturalness theory take it that the most ‘special’ properties – those that do the important work – are the (perfectly) natural ones. These sparse properties are to be distinguished from the abundance of other properties that map on to any grouping we choose to make. Canonical examples of perfectly natural properties include *mass*, *spin*, and *charge*.

Since naturalness itself has such an important role to play in the eyes of its defenders, we should expect that naturalness is itself highly (if not perfectly) natural. The natural properties are thought to cleave to nature’s joints, and so we should expect the property of being such a property to itself be one of the elite, joint-carving properties. Sider (2011: 139) argues that if naturalness were not natural, all the applications of naturalness would be undermined. My main argument in this chapter is that naturalness isn’t natural. If (as I will argue in section 4) the notion of a natural property does not cleave to an objective joint in nature, then, by the naturalness theorist’s own lights, taking the distinction between the (perfectly) natural properties and the rest to be metaphysically significant is unwarranted. It follows that no interesting notion of fundamentality ought to be defined in terms of naturalness.

Lewis (1986b: 61) holds that the distinction between the perfectly natural properties and the rest admits of degree. Properties can be ranked according to their relative degrees of naturalness, with the perfectly natural properties at the limit. There are two ways to characterise naturalness, depending on whether absolute naturalness or relative naturalness is the more basic notion. From each characterisation we can generate the notion of perfect naturalness.⁹

⁹ Lewis (1986: 60) considers the possibility that neither notion is primitive, and that naturalness ought instead to be analysed in terms of some underlying posit. If this is the case, similar considerations to those addressed here ought to apply to that underlying posit, and so for simplicity I will assume here that some notion of naturalness is primitive.

If absolute naturalness is basic, entities can be described with the one place predicate ‘ P is perfectly natural’ (Taylor, 1993: 84). To establish that some property P_1 is more natural than P_2 , we define both P_1 and P_2 in terms of the perfectly natural properties, and compare the definitions. Lewis’ suggestion (1984: 228; 1986b: 61) is that P_1 is more natural than P_2 if the definition of P_1 in terms of the perfectly natural properties is less complex than that of P_2 . This is commonly understood (e.g. Sider 2011: 130) as a matter broadly of the length of the definition of the property in a language where all predicates stand for perfectly natural properties and relations (making allowances for, for example, taking disjunctions to be indicative of greater complexity than conjunctions).¹⁰

If relative naturalness is the more basic notion (and the two-place predicate ‘ P_1 is more natural than P_2 ’ is primitive) then a perfectly natural property is such that no property is more natural than it. This account (unlike the absolute conception) is consistent with relative naturalness being only a partial ordering (where there might be areas of incommensurability across chains of properties with increasing degrees of naturalness, and perfectly natural properties at the end of each chain).

An assessment of whether the property *being perfectly natural*, or *perfect naturalness* fulfils some central components of the role of a perfectly natural property reveals that *perfect naturalness* cannot be perfectly natural. Attempts to assess the relative degree of naturalness had by the property of perfect naturalness further reveal that *perfect naturalness* is in fact a fairly unnatural property.

2. The naturalness role

Recall that naturalness is usually taken to be primitive, in the sense that no full explanation of naturalness can be given in terms of anything else. Instead, naturalness is often introduced by example: the property *being green* is thought to be more natural than the property *being grue*;¹¹ the property *being a plant* is thought to be more natural than the property *being one side of a red die or Barack Obama’s left hand*. Naturalness is a theoretical posit, and like other such posits, we can take its meaning to be fixed by a

¹⁰ One reservation: Lewis says little about how this relative complexity is to be understood. In the absence of a clear account of simplicity, we are given no reason to think that an analysis of naturalness in terms of simplicity will be successful.

¹¹ The predicate ‘grue’ applies to all things observed before time t just in case they are green, and to other things just in case they are blue (see Goodman, 1954).

collection of its roles or theoretical uses. My aim in this chapter is to ask whether *perfect naturalness* itself has the features that count towards a property's being perfectly natural according to those roles. In what follows I present these roles as necessary conditions on a property's being perfectly natural. This comes with a caveat: not all friends of naturalness agree that each of these roles must be satisfied by every perfectly natural property, and many also consider some of the roles to be greater indicators of perfect naturalness than the rest. The roles I consider in this section are: empirical discovery and figuring in laws; providing a minimal supervenience base; intrinsicality and duplication; and reference magnetism.

There are a few other indicators of perfect naturalness mentioned in Lewis' works. For example, he claims that the perfectly natural properties are simple, that they are highly specific, and that the naturalness facts are non-contingent (see e.g. Lewis 1983; 1986b: 60). I'll briefly say something about each of these indicators, before considering the other aspects of the naturalness role in detail. Non-contingency of the naturalness facts is an assumption I make here for the sake of simplicity. I hope this assumption can be considered fairly benign, since it is both shared by Lewis and most of his followers, and does not affect most of the arguments made here. Those that it does affect (e.g. in section 2.2) can be reformulated in terms that do not assume non-contingency.

The idea that the more natural a property is, the more simple it is is understood as the claim that one property is more natural than another iff the former has a definition in terms of perfectly natural properties that is simpler than any such definition of the latter. This is discussed with regards to *perfect naturalness* in section 3. Related to simplicity is the claim that perfectly natural properties are highly specific. We can note immediately that the property *being perfectly natural* is not highly specific, because there are a number of ways a property might be perfectly natural (corresponding to the different perfectly natural properties). *Charge* is perfectly natural in virtue of being *charge*, *spin* in virtue of being *spin*, and so on. The discussion below concerns aspects of the naturalness role that *perfect naturalness* less obviously fails to satisfy.

2.1. Empiricism and laws

It is an aspect of the naturalness role that P is a perfectly natural property iff P appears in fundamental physics. Lewis is clear that if there is an inequality amongst properties (so that a minority are to be considered elite) it ought to be the fundamental physical properties as discovered and named by physicists that have the elite status (e.g. Lewis 1984: 228; 1986: 60). Modern day physicists talk in terms of properties that modern day metaphysicians class as natural properties (e.g. *charge*, *mass*, *spin*) but there is no talk of a property of *naturalness* amongst physicists, and no reason to think that this is likely to change. By the physicists' lights, their task of giving a characterisation of fundamental reality is complete once an inventory of the fundamental entities and their interactions (however the physicists understand them) has been given. There is no need for physicists to add that the properties and relations they have discovered are also perfectly natural.

A closely related condition on naturalness is that P is a perfectly natural property iff P figures in the best deductive system(s). Lewis takes the laws of nature to be the axioms of whichever theory best maximises simplicity and strength (where strength is a matter of virtues such as predictive power, unity, coherence, and consistency). Just as science discovers the perfectly natural properties, it also discovers the laws, and so the two are inseparable (Lewis, 1983: 365). It is the perfectly natural properties that will feature in the laws of nature, where those laws are the laws of the best deductive system(s). But whilst the laws of nature at this world seem to involve perfectly natural properties like *charge*, *mass*, and *spin*, there is no mention of *perfect naturalness* by scientists, and so no scientific laws involving *perfect naturalness*.

Perhaps what matters is that *perfect naturalness* plays a fundamental role in our best *total* theory, which might contain some elements not of particular interest to physics (or to the other sciences). Then naturalness might feature in a best total theory, and thus appear in the laws of the best system, irrespective of its apparent absence from empirical investigation. On such an account, *perfect naturalness* still fails to satisfy the element of the role of the natural properties that has them feature in physics, but its claim to naturalness may nevertheless be justified by appeal to more theoretical metaphysical considerations. Such considerations are addressed in the following subsections.

2.2. Supervenience

The perfectly natural properties are taken to provide a complete characterisation of the world, such that all facts supervene on facts about which things have which perfectly natural properties. In Lewis' terms, one of the primary roles of the perfectly natural properties is that of providing a supervenience base. Lewis' taste for parsimony leads him further to take the perfectly natural properties to form a *minimal* supervenience base – 'there are only just enough of them to characterise things completely and without redundancy' (Lewis 1986: 60). It is precisely this kind of claim that recommends naturalness as a definition of fundamentality. We can say then that this aspect of the naturalness role dictates that P is perfectly natural iff P is a member of a minimal supervenience base (MSB).

A minimal supervenience base must fulfil the following two conditions. First, it must be a set S of properties and relations such that that there can be no difference between any two worlds without a difference with respect to at least one of the members of S . This guarantees that everything supervenes on the properties within the base, and thus that the perfectly natural properties characterise things completely. Second, it must be such that there is no subset S' of S that fulfils the first condition. This guarantees that the base be minimal, because should S' fulfil the first condition, the properties that are members of S but not of S' are redundant; a complete characterisation of things can be given without them.

Standard definitions of supervenience are unrestricted – they quantify over all worlds whatsoever. When the A-properties supervene on the B-properties there can be no worlds at which there is a difference in A-properties without a corresponding difference in B-properties. Thus, if all properties supervene on the perfectly natural properties, then membership of the MSB will be non-contingent in the sense that no property can be perfectly natural *at any world* without being a member of the base. This allows for the possibility of alien natural properties (i.e. properties that do not exist at this world, but exist and are perfectly natural at some other(s)), but does not allow that properties be perfectly natural at some worlds at which they exist, but not at other worlds at which they exist.

It might be that there are multiple candidate sets of properties that can fulfil the conditions for being an MSB. For example, we might replace any member of the set S of perfectly natural properties that make up the MSB with its negation (e.g. we can replace *having mass* with *not having mass*), thus generating an alternative MSB that still fulfils the conditions of supervenience and minimality. Even if we were unable to specify a unique MSB, we could still consider whether the property of *perfect naturalness* itself might be a member of *some* MSB. We turn to this task now.

The Lewisian line is that when two possible worlds differ about the truth value of at least one proposition, they differ with regards to the truth value of at least one proposition that predicates a perfectly natural property of some entity (or a perfectly natural relation of some entities) (Dorr and Hawthorne, 2013: 7). Because no MSB contains any properties that supervene on other properties within the base, for every property within the MSB there ought to be at least one proposition whose truth value differs only depending on that property's being predicated of some entity. This guarantees that every property which is a member of the set S of properties that make up the MSB earns its place, and therefore that the base is truly minimal. (Of course, some members of S might be responsible for a difference between a pair of worlds W_1 and W_2 neither of which is the actual world, but this is just another way of saying that there might be alien perfectly natural properties).

The only propositions whose truth value will differ only depending on the predication of 'is a perfectly natural property' will be propositions *about* which properties are perfectly natural. Orthodoxy has it that which properties are perfectly natural is not a contingent matter (see e.g. Lewis 1986b: 61) and so any proposition predicating perfect naturalness of some property will be necessary. There is therefore no proposition whose truth value differs based on the predication of 'is a perfectly natural property' of some property, and so according to any specification of MSB, *perfect naturalness* will not be a member of S because that property fails to satisfy the requirement that each member of S should be responsible for a difference between some two worlds.¹²

¹² If the naturalness facts are contingent in the sense that a property can be perfectly natural at some world in which it is instantiated, and not at others, then propositions about which properties are perfectly natural will have different truth values in different worlds, and so it is not straightforwardly the case that a predication of 'is a perfectly natural property' cannot be

Given that no set of properties of which perfect naturalness is a member can form a *minimal* supervenience base, a defender of the supervenience aspect of naturalness theory might suggest weakening the characterisation of a supervenience base so that minimality is no longer central, in the hope that *perfect naturalness* might then prove a legitimate member of the base. This could be done in one of two ways: (a) by allowing ‘near-enough’ natural properties to feature in the base, or (b) by specifying conditions under which properties that appear to supervene on one-another might nevertheless both be included in the base. I will argue that any such weakening that retains the spirit of the supervenience aspect of the naturalness role will not allow *perfect naturalness* to feature in the set of properties that comprise the base.

The first argument for weakening the minimality condition on the supervenience base comes from the apparent possibility of infinite complexity, both in the actual world and in other possible worlds. There are various plausible scenarios involving certain kinds of infinite complexity that make it the case that properties are endlessly supervenient on lower-level properties (Schaffer, 2004: 99). For example, it might be that the world is ‘gunky’ – that matter is infinitely divisible.¹³ Just as proton-properties supervene on quark-properties, it might be that the quark-properties supervene on lower-level properties corresponding to smaller sub-atomic particles that compose the quarks, and so on ‘all the way down’. At such worlds, there can be no set of properties that characterise things completely and without redundancy. For any property *P* that is a candidate for inclusion in the set of perfectly natural properties, there are some other properties upon which facts about *P* supervene, thus rendering *P* redundant in any characterisation of an MSB.

responsible for a difference in the truth value of a relevant proposition. However, when we have a conception according to which different properties may play the naturalness role at different worlds then the truth value of propositions predicating perfect naturalness of some property can be accounted for in terms of whichever features of the property of which perfect naturalness is predicated. There is then no need to include perfect naturalness in a MSB.

In worlds where we have no handle on perfect naturalness via other routes (i.e. no idea about the features of the perfectly natural properties) we will need to include perfect naturalness in a MSB in order to give an account of which properties are perfectly natural. In such worlds, however, we will have no positive conception of the naturalness role. If *being a perfectly natural property* is the kind of property we have no positive conception of, it is hard to see how perfect naturalness could be taken to play any significant metaphysical role in the first place.

¹³ The possibility of gunk makes trouble for more than just the conception of fundamentality in terms of naturalness – see Chapter 3.

Lewis himself takes infinite complexity to be a genuine (albeit far-fetched) possibility, and suggests that we might therefore take the ‘near-enough’ natural (NE-natural) properties to play the natural property role (Lewis, 2009: 218). Lewis holds that these NE-natural properties can form a base (though not, of course, a minimal base) upon which everything else supervenes. Whatever properties are the NE-natural ones, the property of *being an NE-natural property* (which will be the relevant property if NE-natural properties play the naturalness role), will supervene on the NE-natural properties themselves. *NE-naturalness* will therefore not form part of the supervenience base.

An alternative way to deal with the relevant kind of infinite complexity whilst continuing to endorse naturalness theory is to reject what Schaffer (2004) calls the ‘fundamental’ conception of perfectly natural properties, opting instead for a ‘scientific’ conception, where perfectly natural properties are drawn from every level of a scientific characterisation of the world. On such a view, economic and psychological properties, along with physical and chemical properties might all contribute to the set of perfectly natural properties. What is important is that the perfectly natural properties provide an ontological base that makes true all of the truths at a world (Schaffer, 2004: 100).

It is not clear that even on this sort of account, *perfect naturalness* ought to count as a perfectly natural property. There is little reason to think that *perfect naturalness* enters into a scientific characterisation of the world at any level, and less still to think it plays a role in making true any of the truths at a world that are not already made true by the perfectly natural properties. Fundamental metaphysics is not a special science.

The second way in which the characterisation of the base might be weakened (in such a way that the base might include *perfect naturalness*) does not turn on infinite complexity. Fine (2001: 11) offers the *mass*, *volume* and *density* of a homogenous fluid as a case of three parameters where the value of any one parameter supervenes on the values of the two others. We cannot, on pain of arbitrariness, stipulate which two parameters ought to be included in an MSB, and for each parameter there is good reason not to include it – that its value supervenes on the values of the other two. Nevertheless, unless at least two parameters are included, we will be unable to give a complete characterisation of things. It is therefore

impossible, if *mass*, *density* and *volume* are taken to be perfectly natural properties, to give a complete, non-arbitrary and non-redundant characterisation of things.

A defender of naturalness theory might attempt to use this example to show that there can be members of an MSB that supervene on other properties within the base. But this is really just a counterexample to the idea that there is a single, *unique* MSB, and the idea that there is a *unique* MSB is one we have already been forced to forgo (we saw above that we could generate alternative MSBs by replacing one or more of the properties in the set with its negation). The example above indicates that there are at least three minimal supervenience bases, each of which contains a different pair of the three parameters discussed. Our question is whether *perfect naturalness* features in *any* MSB, and it seems it does not.

Whilst any two of *mass*, *density* and *volume* might appear in an MSB, no MSB will include *perfect naturalness* because *perfect naturalness* supervenes on all of the properties within the base (since, by hypothesis, the properties that form the MSB are all and only the perfectly natural ones). Any set *S* of properties that has *perfect naturalness* as a member will have a proper subset that excludes *perfect naturalness* and that nevertheless characterises things completely. The property *perfect naturalness* is always redundant in giving a complete characterisation of things, whereas each of *mass*, *density* and *volume* can only be considered redundant if the other two parameters are members of the MSB. The characterisation of an MSB cannot be weakened in such a way that the MSB might include *perfect naturalness* and still play the required aspect of the naturalness role.

There is a further, related aspect of Lewis' characterisation of the perfectly natural properties that prevents *perfect naturalness* from fulfilling the supervenience aspect of the naturalness role. Lewis (e.g. 1983: 358) argues that the perfectly natural properties ought to be mutually independent. Even on the weakest formulation of mutual independence, which holds that facts about no perfectly natural property supervene on facts about all of the others, *perfect naturalness* cannot be considered a perfectly natural

property; facts about *perfect naturalness* supervene on facts about all of the other perfectly natural properties.¹⁴ This argument is independent of any considerations about a minimal supervenience base.

2.3. Intrinsicity and Duplication

For Lewis, intrinsicity and duplication are related in a tight circle of interdefinition. All perfectly natural properties are intrinsic, though not all intrinsic properties are natural (Lewis 1983: 357). (The property *being person-shaped* is a plausible example of an intrinsic, but not particularly natural, property; whilst *having a positive charge* is perhaps both intrinsic and perfectly natural.) A property *P* is perfectly natural iff it makes for perfect intrinsic duplication. Perfect duplicates share all of their perfectly natural properties, and the parts of perfect duplicates can be put into correspondence in such a way that corresponding parts have exactly the same perfectly natural properties (and stand in exactly the same perfectly natural relations) (Lewis, 1986b: 62). Two things are perfect duplicates iff they have exactly the same intrinsic properties, and intrinsic properties never divide duplicates within or between worlds (Langton and Lewis, 1998: 336).

Imagine a machine that creates perfect duplicates in accordance with Lewis' recipe. If the machine duplicates electron *e*, it must take into account factors such as the *mass*, *spin* and *charge* of *e* – those are its natural, and therefore intrinsic properties. The machine need not take into account *e*'s location, which objects it is a part of, or how it moves relative to external observers – those are extrinsic and therefore not perfectly natural properties. *Perfect naturalness* is a property of properties, and we can assume that the naturalness of a property is intrinsic to it.¹⁵ If *perfect naturalness* were itself perfectly natural then when the perfect duplicating machine duplicated a given first order property, it ought to take into account the degree of naturalness enjoyed by the property being duplicated. If we could show that property

¹⁴ At least, if it is a non-contingent matter which properties are perfectly natural. If properties can be perfectly natural in some worlds at which they are instantiated and not in others, then facts about which properties are perfectly natural won't supervene on all the facts about which properties happen to be perfectly natural at the actual world. But it is the case that facts about which properties are perfectly natural will be determined by properties about the properties that *are* perfectly natural at a given world. Even if the naturalness facts are non-contingent, *perfect naturalness* and the other perfectly natural properties are not mutually independent.

¹⁵ It is notoriously difficult to define the notion of intrinsicity, and to do so with appeal to naturalness would be question begging in this context. Nevertheless, the extent to which a given property is natural seems to be given purely by the way the property is; not in virtue of the way in which that property interacts with the world. The degree to which a given property is natural seems not to 'imply accompaniment' (see Kim, 1982; Langton and Lewis, 1998: 333).

duplication involved such a consideration, it would suggest that *perfect naturalness* fulfils the duplication aspect of the naturalness role. But property duplication is immediately problematic. To see this, we can consider how various theories about the nature of properties might account for property duplication.

First, consider the view that properties are abstract universals. Discounting (because it offers up no answers to the question at hand) the trivial sense in which everything might be said to be a duplicate of itself, if there were duplicate universals it would need to be possible for there to be two of the same universal. This would undermine much of the motivation for appealing to universals in the first place. Here's an example: we appeal to the universal *green* in order to give an account of what this green chair and that green ball have in common. If there are multiple *green* universals, there is no guarantee that any universal is shared by the chair and the ball (perhaps there is a different universal corresponding to each green thing).

The above problem will also apply if properties are concrete, immanent universals (such that the universal is multiply located in its instances); there is no universal that is guaranteed to be shared by the different objects that instantiate the property in question. A further problem also applies if properties are immanent universals. On this conception of properties, we cannot feed *being green* into the perfect duplication machine, because we cannot extract *being green* from its instances. Suppose then we feed a green apple into the perfect duplication machine. The machine will duplicate all of the perfectly natural features of the apple, and the duplicate will indeed be green. But all we have done is duplicated the apple. The greenness of the apple was duplicated in virtue of its supervenience on the perfectly natural properties of the apple, and thus the machine had no need to take the naturalness status of *being green* into account.

Suppose instead that we take properties to be sets of their instances. Under this conception, to duplicate a property would be to somehow duplicate the set of all of its instances. Since both the original

and the duplicate sets would have the same extension, this is a violation of the extensionality axiom of set theory, which states that sets with the same elements are identical (i.e. are the same set).¹⁶

Perfectly natural properties are shared between duplicates, but property duplication seems metaphysically impossible. Consequently, the duplication aspect of the naturalness role is at best only trivially satisfied when the entity to be duplicated is itself a property. Whilst *being perfectly natural* doesn't violate the relevant condition on perfectly natural properties (that they not divide perfect intrinsic duplicates within or between worlds), non-violation of that condition gives us no reason to think that perfect naturalness is perfectly natural. As applied to properties, this constraint on perfectly natural properties does not give us an informative criterion for assessing perfect naturalness.

Worse, like *perfect naturalness*, properties such as *being a property* will also trivially satisfy the duplication role of a perfectly natural property. If we grant that the criterion is trivially satisfied with respect to *being perfectly natural*, we must also accept that it gets what is intuitively the wrong result when applied to other second order properties. The false positives that are generated by this trivial satisfaction of the duplication role at the level of second order properties casts some more general doubts on Lewis' project of outlining the naturalness role. A Lewisian conception of naturalness builds a number of specific commitments in to the theory (for example, the idea that for a property to be perfectly natural, it must be metaphysically possible to duplicate whatever instantiates that property). This should be considered a cost for naturalness theory, because it makes its acceptance contingent on the acceptance of other metaphysical baggage.

2.4. Reference magnetism

Lewis (1984: 227) argues that *ceteris paribus*, the natural properties provide more *eligible* referents for the predicates of our language. They are the properties that are most easily picked out by our expressions; the 'reference magnets' (see appendix for a more detailed discussion of reference magnetism than I can provide here). A property *P* is perfectly natural iff *P* is a *reference magnet*. On the usual

¹⁶ Note that there is plausibly a subtle difference between duplicate sets formed by duplicating their members, and duplicating sets themselves. An example of the former case would be taking the unit sets of two duplicate electrons to be duplicate sets. Here the sets can be distinguished through differences in the extrinsic properties of their members. In the latter case of duplicating the sets themselves (which I claim here is impossible), the sets could not be so distinguished.

interpretation (see e.g. Sider, 2009; 2011; Weatherson, 2003) Lewisian metasemantics involves constructing an interpretation of a language that best balances two factors. The first is use, where we are to interpret people's actions, and therefore their mental states (and, indirectly, their language) as maximally rational, given their environment. The second is eligibility, where eligibility is accounted for in terms of naturalness. Crudely, the thought is along the lines that even if our theory about electrons contains some false propositions about, for example, *charge*, *charge* is so eligible to count as the referent of our expression 'charge' that it is still *charge* we refer to, even though we believe and state some false things about it. (The extreme alternative would be to hold that since we intend to speak truly, the referent of our expression 'charge' must be whatever property makes all of our platitudes concerning charge come out as true, presumably a highly gerrymandered and unnatural property).

Our question here, then, concerns the eligibility of *perfect naturalness* to serve as the referent of our expression 'perfectly natural'. But we as speakers of the language are in no position to make judgments concerning the eligibility of particular things to be referents. The alleged strength of reference magnetism depends on the eligibility of things to be referents being an *external* constraint on the metasemantic theory (see Lewis, 1984). Our intention to refer in such a way that our platitudes come out as true will be the same whether or not any of the candidate referents for our expression are perfectly natural, and so we will be unable to judge which are the reference magnets based on our own attempts to refer (see appendix for a more detailed exposition of this argument). We can't then hope to discover how eligible *perfect naturalness* is to be a referent of our expression 'perfectly natural'.

A closely related question may shed light on the issue. Dorr and Hawthorne (2013: 52-60) are interested in assessing the level of *vagueness* in 'perfectly natural'. If the expression 'perfectly natural' is vague, that would be a good indication that *perfect naturalness* is not very natural. If, on the other hand, 'perfectly natural' is precise, one good explanation (assuming the truth of reference magnetism) for its precise nature is that there is a joint carving interpretation of the expression.

There are (at least) two ways that 'perfectly natural' might be vague. First, there might be multiple equally good candidate interpretations of the expression (i.e. there might be multiple sets of properties that each fully satisfy the naturalness role, with none eligible enough to fix a determinate referent). This

would suggest that the property *perfect naturalness* is not a reference magnet, and thus suggests that *perfect naturalness* is not perfectly natural. Second, there might be *no* set of properties that definitely satisfies the naturalness role (even though the role is fully satisfied by the perfectly natural properties). In this case, there is no set of properties that are all and only the perfectly natural ones (see Dorr and Hawthorne, 2013: 53). The extension of ‘perfectly natural’ is vague because no set of properties is picked out by it.

One reason given by Dorr and Hawthorne (2013: 55) for thinking that ‘perfectly natural’ is at least somewhat vague is that there are some interdefinable properties (e.g. Fine’s (2001) example of *mass*, *density* and *volume* in a homogenous fluid; some properties of Euclidean geometry; some mereological properties) such that one ought to be taken as perfectly natural and the other defined in terms of it, but there is no reason to prefer taking either one as perfectly natural over the other. An attractive option is to claim that there is no determinately correct answer to the question of which of these properties are the perfectly natural ones. There is no reason to prefer, for example, *parthood* over *proper parthood*, or vice versa, as the more natural notion.

Perhaps each of these properties is included in a set that is one of multiple equally good interpretations of the expression ‘perfectly natural’. In this case, the set that includes *parthood* does not also include *proper parthood*, and so on for other interdefinable properties. The expression ‘perfectly natural’ would then be vague in the first of the two ways specified above. Lewis’ metasemantic theory takes use seriously to constrain the interpretation of a language. If use of the term ‘perfectly natural’ privileges minimality (as it seems to), we have good reason to think that the examples given above point to there being multiple best candidate interpretations of ‘perfectly natural’, each of which takes a different set of properties that form a minimal supervenience base to be the referent of ‘perfectly natural’.

The alternative (corresponding to the second of the two ways in which ‘perfectly natural’ might be vague) is that the case of properties like *parthood* and *proper parthood* serves to indicate that there is no set (rather than multiple sets) of all and only the perfectly natural properties. It is of course open to anybody who takes ‘perfectly natural’ to be precise to claim that in the above examples one of the properties or relations in question simply *is* the perfectly natural one (or is more natural than the other), and is a member of the unique set of perfectly natural properties regardless of our ability to determine that this is

the case. Though this claim is consistent with naturalness theory, it is also costly. It is generally a bad idea to believe things we have no clear evidence for.

A different reason to think that there might be a reasonable degree of vagueness in the expression ‘perfectly natural’ is that *perfect naturalness* is taken to be a theoretical primitive, and thus is not given an explicit definition. *Perfect naturalness* is introduced by example and by explicating the theoretical role it is taken to play, but one might worry that this will not be enough to eliminate vagueness from the expression ‘perfectly natural’ (compare the sceptical challenge mounted against the grounding relation in Daly, 2012 (and discussed in Chapter 1, section 2.2.1)). As in the examples of vagueness given above, ‘perfectly natural’ might be vague even when the entire naturalness role is satisfied.

It seems likely that the expression ‘perfectly natural’ is at least moderately vague. This might be independently worrying for the naturalness theorist (it is presumably better to have precise theoretical posits than to have vague ones). Moreover, it is unlikely that *perfect naturalness* is an extremely eligible referent for the expression ‘perfectly natural’ if *perfect naturalness* is moderately vague. This suggests that *perfect naturalness* is not a reference magnet, and therefore fails to fulfil yet another of the roles that the perfectly natural properties are taken to play.

3. How natural is perfect naturalness?

The arguments of the previous section show that *perfect naturalness* fails to fulfil some key aspects of the naturalness role. It is not the kind of thing that is discovered or identified through empirical investigation or is involved in physical laws; it fails to form part of a minimal supervenience base; and there is little sense to be made of it playing a role in duplication, or being intrinsic to its bearers. The apparent degree to which ‘perfectly natural’ is vague suggests that *perfect naturalness* is not a reference magnet. Putting this all together suggests that *perfect naturalness* is not a perfectly natural property.

Whilst from our assessment above it seems that *perfect naturalness* is in fact quite unnatural (because of the ways in which it falls short of fulfilment of various aspects of the naturalness role) it is not obviously legitimate to infer from the extent to which *perfect naturalness* fails to play roles of a perfectly natural property, that it is anything other than not perfectly natural. We might then look to an alternative

method for assessing the degree to which *perfect naturalness* is natural. As noted above, the Lewisian method for assessing the relative naturalness of a given property P is to consider the complexity of P 's definition in terms of the perfectly natural properties (e.g. Lewis, 1986b: 61). A definition of a property in this context is an expression which provides necessary and sufficient conditions, and definitions in terms of perfectly natural properties will be expressions in a language where all syntactically simple non-logical vocabulary expresses perfectly natural properties (Dorr and Hawthorne, 2013: 15).

Sider (2011: 138) proposes that if *perfect naturalness* itself is not perfectly natural then the following is a definition of *perfect naturalness* (where N_n are all and only the perfectly natural properties):

P is perfectly natural =_{df} $P=N_1$ or $P=N_2$ or... $P=N_i$

This definition is highly disjunctive, which makes it appear highly unnatural.¹⁷ On this method of assessing naturalness, if *perfect naturalness* is not perfectly natural, then it is highly unnatural (because it is defined in the above way).

Here's one way the friend of naturalness might respond: *This can't be the right way to define perfect naturalness. The N terms are functioning in the definition as names for perfectly natural properties, rather than as predicates. If we allow names to feature in the definition in this way, then some intuitively very unnatural properties will turn out to be extremely natural.*

But as Dorr and Hawthorne, (2013: 16, note 30) explain, if we allow names for properties and a predicate 'instantiates', we can define any property extremely simply with 'instantiates P ', where P is a name for the property in question. We would then have to count all properties as equally natural, because they all have a definition of equal length.

So we have good reason to think that names should be excluded from all definitions of properties in perfectly natural terms. How then are we to define *perfect naturalness*? One suggestion might be to allow only a limited class of property names – those that denote the perfectly natural properties – to feature in the relevant definitions. But this will not help the friend of naturalness, because in that case the

¹⁷ Exactly *how* complex and disjunctive the definition is depends on how many perfectly natural properties there are. Given that there must be enough to characterise things completely, it is unlikely that the perfectly natural properties will be sparse enough to maintain that the disjunction of all of them is simple enough to indicate a very natural property.

definition of *perfect naturalness* will be the complex, disjunctive definition given above. In the absence of an alternative method for assessing how natural non-perfectly natural properties are, the friend of naturalness is unable to give an account of how natural *perfect naturalness* is. This is something of a dilemma for the friend of naturalness: she must either accept the complex, disjunctive definition given above; or she must come up with an alternative method for assessing the naturalness of less-than-perfectly natural properties (and show that by its lights, *perfect naturalness* is at least fairly natural).

The first horn of the dilemma is the conclusion we can assume the friend of naturalness is trying to avoid (its consequences are discussed in the next section). Though the second horn might yet prove fruitful, there is much work to be done even to come up with a systematic method for assessing naturalness, let alone to demonstrate that *perfect naturalness* will turn out highly natural by its lights.¹⁸ The only other thing the friend of naturalness can do is to try to reject the dilemma altogether. I turn to that response now.

Perhaps there is an implicit restriction within naturalness theory such that it applies only to first-order properties (i.e. properties of non-properties). There are a few ways such a restriction could be cashed out: perhaps it is only first-order properties that have a definition in perfectly natural terms; perhaps the naturalness roles are to be filled only by first-order properties; or perhaps only first-order properties can be ranked according to their degrees of relative naturalness. In any of these cases, questions about the naturalness of perfect naturalness, a second-order property, would be illegitimate, or confused, and so the dilemma presented above would be a false one.

There are two replies to the suggestion that naturalness theory might be restricted in one (or more) of these ways. The first is that the claim has an air of ad-hocery about it. Questions about the naturalness of *perfectly natural* certainly seem grammatical and well-formed in English (we have, for example, spent the last while attempting to answer them as if they were well-formed). In any case, it is not always a mistake to apply predicates to higher-order properties. The predicate ‘is a property’, for example, seems as

¹⁸ One possibility we might consider is looking at how well a given property *P* fulfils the roles of a perfectly natural property, to give us an indication of how natural *P* is. Those roles are fairly well defined, and we can perhaps get a reasonably good handle on the extent to which *P* falls short of them. For *perfect naturalness* though, this assessment would suggest a fairly unnatural property (see section 3).

though it can be correctly applied to a property of any order. To make predicates involving naturalness a special case certainly adds to the complexity of the theory, and is thus an unanticipated theoretical cost for the defender of naturalness who embarks down this road in response to the problems identified above.

The second reply appeals to an argument for the necessity of perfectly natural second order relations discussed in Eddon (2013). Eddon's argument is that in order to give an account of quantitative properties (e.g. *two grams mass*) that preserves a necessary ordering of such properties (e.g. *two grams mass* is less than *three grams mass*) and distance between them (e.g. the distance between *two grams mass* and *three grams mass* is one gram), we require perfectly natural second-order relations holding between the quantitative properties.

Without such relations, there is no plausible way (see Eddon, 2013 for a discussion and rejection of possible alternatives) to give an account of why it is that, for example, two objects instantiating *2 grams mass* and *3 grams mass* respectively resemble one another more closely than they do an object instantiating *1000 tonnes mass*. Eddon argues that appeal to first order natural properties cannot distinguish between them since each object instantiates different quantities of mass, and thus none of them share any perfectly natural properties (appeal to the determinable *mass* that each instantiates fails to capture degrees of resemblance between them). The difference cannot, Eddon argues, be explained by appeal to non-fundamental second-order relations since there will be myriad such relations and no way to distinguish any which is privileged. Making the appropriate distinctions requires appeal to perfectly natural second-order relations.

Eddon thus provides good reason to think that there must be perfectly natural second order relations.¹⁹ If there are some such relations, it becomes highly implausible that *perfect naturalness* is restricted so as to range only over first-order properties and relations. To claim that the restriction on perfect naturalness is

¹⁹ This has the consequence that the minimality criterion in the account of the role that the perfectly natural properties play in comprising a supervenience base must be rejected. Any perfectly natural second order properties violate this aspect of the role in the same way as *perfect naturalness* does if we take it to be perfectly natural (though not to the same extent – perfectly natural second order relations of this kind do not supervene on *all* the other properties within the MSB). However, the relations that Eddon discusses do play most of the other aspects of the naturalness role (e.g. in duplication, laws, evaluating counterfactuals, etc.) which is her motivation for taking them to be perfectly natural in the first place (see Eddon, 2013, section 5). The presence of some perfectly natural second order relations does not pose a threat to my overall argument here.

only such that *perfect naturalness* itself cannot be evaluated in terms of naturalness is ad-hoc and therefore cannot be maintained. *Perfect naturalness* is an unnatural second order property.

4. Consequences

What are the consequences if, as the arguments of the previous section give us reason to believe, naturalness is highly unnatural? It might be tempting to think that if *perfect naturalness* is not a natural property, then the perfectly natural properties are not natural either. Without further argument, this is just a confusion. Suppose some property *P* is a simple property, but the property *being a simple property* is extremely complex (highly disjunctive, for example). The complexity of *being a simple property* need not infect the simplicity of *P*. It would be obviously mistaken to infer from the complexity of the property of simplicity that properties with the property of simplicity are thereby complex, and it is similarly mistaken to infer from the unnaturalness of *perfect naturalness* that perfectly natural properties are in fact not perfectly natural.

The unnaturalness of *perfect naturalness* does not immediately infect the natural properties themselves, but it is independently problematic. This is so for (at least) two reasons: the first reason is methodological (it undermines the explanatory role that the natural properties are taken to play); and the second reason is metaphysical. I'll explain these in turn.

The suggestion that *perfect naturalness* might be fairly unnatural is considered by Sider (2011: 138-9) under the title of 'Melianism' (after conversation with Joseph Melia). Sider's Melian embraces the unnaturalness of *perfect naturalness*, on the grounds that *perfect naturalness* itself is never involved in explanations, only particular natural properties. The claim is that it is not a problem that *perfect naturalness* itself is not included in the fundamental picture of the world, because the perfectly natural properties are. When we give accounts that appeal to naturalness, they in fact appeal to particular (perfectly) natural properties. This is the case even if we are unaware which of the natural properties is playing a particular role – the correct account of things invokes some particular natural property or other, and so there is no need for inclusion of the general notion of naturalness in that account.

Sider (2011: 139) thinks that Melianism undermines all of the applications of naturalness. Sider has a conception of naturalness theory as fundamental in our description of the world because it offers an explanation of important metaphysical aspects of reality, as detailed in section two above. Sider claims that if naturalness is as unnatural as the Melian supposes, naturalness theory could not be explanatory. The Melian responds by citing only particular natural properties to do the necessary work in providing an explanation of a given phenomenon. Sider counters that explanations require generalisations (and they must be generalisations of sufficient scope). For generalisations about naturalness to be explanatory, naturalness cannot have a Melian definition.

Consider an event such as the breaking of a window. We can provide an explanation for that event in terms of properties such as the fragility of the window and the force of the ball that came flying at it. What we cannot do, if naturalness is not natural, is explain why that makes for a good explanation (i.e. because it cites properties that are natural enough in the relevant context, and the explanation depends on some laws of nature which involve perfectly natural properties). As Sider (2011: 139) argues, the unnaturalness of (*perfect*) *naturalness* makes attempts to give accounts of lawhood, duplication, reference etc. in terms of (*perfect*) naturalness seem an arbitrary exercise, and one not especially worth pursuing.

The second reason that naturalness must be at least fairly natural does not hinge on the explanatory role of the natural properties. The (*perfectly*) natural properties are supposed to be the important ones; those that carve reality at the joints; those that make for objective distinctions in reality. If *being a perfectly natural property* doesn't itself carve reality at the joints and make for an objective distinction, then there is no relevant difference between this and other, seemingly less significant distinctions. The distinction between the perfectly natural properties and the rest becomes like the distinction between properties that are common between things that are green, and properties that are not – it is not the sort of distinction that we ought to care about in metaphysics. Naturalness theory has it that the metaphysically significant properties are the natural ones, and so *perfect naturalness* must be at least fairly highly natural if naturalness is to be metaphysically significant. Given that naturalness theory is premised on the idea that naturalness *is* metaphysically significant, that *perfect naturalness* seems fairly unnatural undermines the enterprise.

If naturalness is not natural, and we understand fundamentality in terms of naturalness, then fundamentality is itself not fundamental. Those who think fundamentality talk gives us a better way to represent the world than we are able to without it would occupy a strange position if they thought the very notion of fundamentality was itself gerrymandered and unnatural, related to the fundamental only by a long chain of definition.

5. Structure

If naturalness isn't natural, then an analysis of fundamentality in terms of perfect naturalness starts to look less promising. The obvious next place to search for an adequate conception of fundamentality is in Sider's (2011) extension of naturalness theory. Recall that, according to Sider, x is fundamental \equiv_{df} x is part of the world's structure (where 'structure' here refers to Sider's primitive theoretical posit).

The problem for naturalness theory that I identified above comes about primarily because naturalness theory seems to break down when it is applied beyond first order properties. This we saw to be explicitly the case in our analysis of the duplication role, but it was lurking in the background of the discussion of supervenience, and to some extent in the discussions of empiricism and laws and reference magnetism too. The problems with giving a definition of perfect naturalness can be diagnosed along the same lines.

Sider's (2011) proposal extends naturalness theory 'beyond the predicate', freeing it from the constraints that seem to generate the problems discussed above. Recall that the primitive structural operator Sider introduces (2011: 92) applies to any arbitrary portions of the language, including to quantificational phrases and even to logical connectives. For anything we might wish to express, Sider's extension of naturalness theory provides a mechanism for us to consider whether or not it is part of the fundamental structure of the world.

If we ask the analogous question to the one we have been considering here of structure (is 'structure' structural?) we should expect Sider to answer 'yes'. (The ways in which the unnaturalness of naturalness is unpalatable has a straightforward extension to structure). The question is easier to ask of Sider than it is of the friend of naturalness, because Sider's posit is designed to attach to any portion of the language

(including operators like ‘structure’). Sider indeed considers this question explicitly, and his initial response is simply ‘yes’; structure is structural (2011: 137).

Sider’s primary reason for embracing the structuralness of structure is the methodological one discussed in the previous section. If structure were not structural, then structure itself would fail to be explanatory. Any definition of structure in structural terms would have to be long and disjunctive, since the class of structural notions is so heterogeneous (it includes, for example, mass, spin, set membership, disjunction, and universal quantification). Nothing unifies the class of structural notions besides their all being structural (Sider, 2011: 140).

We are faced, Sider claims, with a choice: ‘either adopt extreme realism about structure—holding that structure is itself structural—or else give up altogether on explanations that invoke structure, which is tantamount to giving up on structure itself’ (2011: 140). It is vital then for Sider’s proposal that structure is indeed structural.

As Schaffer notes, Sider (2011: 128) is explicit about what his primitive is; ‘structure is absolute: I say “is structural” rather than “is more structural”’. Sider’s structure operator attaches to single expressions rather than relating pairs of expressions, and he has good reason for taking the absolute rather than the relative notion of structure to be primitive. Sider subscribes to a principle he calls *purity*: ‘fundamental truths involve only fundamental notions’ (2011: 106). As we have seen, he also takes structure to be structural. If a relative notion of structure were primitive, then there would be fundamental truths that related the fundamental to the non-fundamental. Fundamental truths would then involve non-fundamental notions, violating purity.²⁰

Because structure is an extension of Lewisian naturalness, we should anticipate that a similar argument to the argument I make against naturalness above could be mounted against structure. A thorough investigation is likely to reveal that structure is not structural. Constraints on space prevent me from undertaking that investigation here, though I’ll briefly mention some reasons for thinking that structure isn’t structural. I’ll appeal in part to an argument made by Schaffer (2014) which purports to show that

²⁰ The threat here is what deRosset (2013) refers to as ‘the collapse’ – facts that relate the fundamental to the nonfundamental drag everything into the realm of the fundamental. I discuss this further in the next chapter.

the role Sider carves out for the notion of structure is played by a relative notion of structure, and not the absolute notion Sider himself takes to be primitive. First, I'll introduce some of Sider's own reasoning for the conclusion that fundamental posits are to be avoided in general. If, as I will argue, Sider's primitive does not in fact play the role carved out for it, then Sider's own ideological parsimony ought to lead him to abandon the notion of structure.

As with naturalness, grounding, and other such theoretical posits, appeal to structure is largely justified by the role that structure plays in metaphysics. Sider admits (2011: 141) that his own very strong preference for simplicity militates against many primitive notions. So strong is this preference of Sider's that he thinks even well-established primitive concepts such as modality and causation fall by the wayside. He gives just one good reason for opposing, for example, modal primitivism: 'ideological economy' (2011: 267). As Sider himself admits, when it comes to structure these scruples 'go out the window' (2011: 141).

Sider responds, in a way he concedes is not entirely persuasive, with the claim that structure cannot be reduced without loss (2011: 141). Of course, defenders of primitive modality and primitive causation can say the very same thing about their own posits, but we can set aside their indignation for now. Instead, I want to argue that the notion of absolute structure does not play the role Sider carves out for his primitive. Given Sider's strong taste for ideological parsimony, this diminishing in the applications of Sider's primitive are strong grounds for questioning commitment to it.

First, we should be aware of a complication of Sider's view. Lewisian naturalness is a second order property of properties that play a certain, quite well defined theoretical role. This makes it reasonably straightforward to assess the extent to which the property *being a perfectly natural property* itself plays that role. Many of the roles that Sider carves out for structure are not roles for particular structural *notions* (in accordance with which we could assess the extent to which structure is itself a structural notion) rather they are roles for Sider's primitive itself. In these cases, we can draw on Jonathan Schaffer's argument in his (2014) review of Sider's book, which purports to show that the roles Sider identifies for his primitive are not in fact played by the absolute notion of structure.

Just as Lewis thinks that it is the perfectly natural properties that figure in laws, Sider thinks that the best theories are cast in perfectly fundamental (i.e. absolutely structural) terms (2011: 23). And just as there is no mention of perfect naturalness by physicists, there is no mention of structure, either. Sider will argue that structure figures in the best total theory, but we have no independent reason to think so, and so are justified in claiming that absolute structure falls short of this first element of the structure role.

Analogous to the role that supervenience plays in Lewis' account of perfect naturalness is Sider's characterisation of joint carving. Sider himself stresses that 'a vivid test for whether a given expression, *E*, carves at the joints is this: did God need to think in *E* terms when creating the world' (2011: 138). Sider goes on to admit that God clearly must think in terms of perfectly structural notions like quantification, mass, and distance, but that having done so God need not also consider whether such notions are structural. Again, this suggests that contrary to what is required of it, absolute structure is not structural.

Schafer's (2014) criticism of Sider is encapsulated by what Schaffer calls *mismatch*: "The roles for structure are "structural enough" and "more structural than", but Sider's primitive is "perfectly structural"" (Schaffer, 2014: 126). As with Lewisian naturalness, many of the roles that structure plays do indeed seem to require a comparative notion of structure. Sider himself acknowledges this (2011: 129), but to emphasise the point I'll give a few examples.

One of Sider's primary roles for structure (and a condition that accounts for how structure is able to play many of the other roles carved out for it) is that theories couched in structural terms make for better explanations. But most of the notions involved in explanations are not absolutely structural. Sider concedes this, claiming that genuineness of explanation requires only *somewhat* structural notions (2011: 141). Not only does this exhibit *mismatch* (genuine explanations must be given in terms that are more structural than some other terms, not in absolutely structural terms) but it also demonstrates another sense in which absolute structure is not absolutely structural. The notion of absolute structure does not itself give a characterisation of genuine explanation, and so absolute structure does not play the role carved out for Sider's primitive. Instead, it is a relative notion of structure that is involved in our theorising about explanations; genuine explanations are couched in terms of *somewhat* structural notions.

At both the first order level (which terms are involved in genuine explanations) and the second order level (what theoretical posit is involved in characterising genuine explanation) it is the comparative notion of structure that does the work.

A comparative notion of structure is also required for Sider's characterisation of substantivity, where a question is nonsubstantive just in case 'their answers turn on which of a range of equally good meanings we choose for words in those questions' (Sider, 2011: 46). When no candidate interpretation of the words in the question carves closer to the joints (i.e. *is more structural than*) any other candidate in the vicinity, our question is nonsubstantive. Sider's positive characterisation of substantivity derives from this; a question is substantive when it is not nonsubstantive. The notion of substantivity does a lot of work in Sider's system (working out whether or not a given metaphysical dispute is substantive is crucial to Sider's overall project) but the notion it depends on is ultimately a comparative one.

A final example comes from the very beginning of the book (2011: 1-2). Sider motivates objectivity about structure with an example involving a universe entirely full of fluid, divided by a central plane. Fluid in one half of the universe is uniformly red, and in the other it is uniformly blue. He asks us to consider a group of people who think of the universe as divided by a different, diagonal plane, and thus corresponding to their predicates, 'rue' and 'bred', which cut across our predicates 'red' and 'blue'. Sider's claim is that whilst this imagined community make no mistake in applying their own concepts, they've *got the wrong concepts*. But (at least according to Sider) the book of the world isn't written with predicates like 'red' and 'blue', and so if Sider is to use his example to pump intuitions for the objectivity of structure, then he must concede that it is the idea that 'blue' and 'red' are *more structural than* 'rue' and 'bred' which is doing the work.

To stave off the challenge presented above, Sider owes an account of how to define a notion of comparative structure out of that of absolute structure, and so to explain how absolute structure is fundamental even though it is a comparative notion of structure that does much of the interesting work. But he admits (2011: 129) to having no such account. The best he can do is to suggest that the solution will be something akin to the Lewisian definitional chain method for making comparative judgements of naturalness (though he finds a number of faults with Lewis' own suggestion) and to gesture at four

elements that will figure in a definition of comparative structure.²¹ Even if Sider were able to solve this problem, there is no guarantee that he would thereby be able to show that absolute structure is absolutely structural, only that it is compatible with absolute structure being primitive that a relative notion of structure might play many of the relevant roles. An application of Sider's own preference for a parsimonious ideology renders acceptance of structure impermissible.

6. Conclusion

In this chapter I have discussed two closely related conceptions of fundamentality – naturalness and structure – and have argued that both notions are problematic. If the perfectly natural properties are fundamental, then we should expect *being a perfectly natural property* to itself be perfectly natural, but I have argued that this is not the case. Similarly for Sider's extension of naturalness theory; absolute structure does not play the role that Sider carves out for his fundamental posit.

Lewisian naturalness and Sider's 'structure' both purport to tell us about the structure of reality, but both are subject to problems when we consider whether the posit itself plays the role carved out for the notion with which the structure of reality is to be characterised. The last major contemporary contender for characterising reality's structure is grounding. Though an analogous question to that which has been the source of the problems discussed here can be posed for grounding (namely: is grounding grounded?), grounding is theoretically less constrained than either naturalness or structure, and so has more resources available with which to respond to that question. The difficulties identified for naturalness and 'structure' do not obviously generalise. Nevertheless, there are many outstanding issues with which a theory of grounding must contend. Some of these are the focus of the rest of the thesis.

²¹ Since these elements are poorly defined and, in any case, do not by themselves constitute a definition, I do not review them here.

Chapter 3: Metaphysical Interdependence

In this chapter I turn to discussing grounding, which is the notion most often associated with fundamentality and with characterising the structure of reality in the contemporary literature. A very common understanding of fundamentality is that x is fundamental \equiv_{df} nothing grounds x . The following three chapters are about different aspects of the notion of grounding: its logical properties and their role in theories of metaphysical structure, its connection to the notion of explanation, and the possibility of antirealism about grounding.

Though discussion of grounding has become widespread, much of the literature about grounding is devoted to arriving at a proper characterisation of the notion. Difficulties arise because grounding is taken to be a metaphysical primitive – it resists reductive analysis. This leaves the notion vulnerable to the charge that it has no distinctive content, or no useful role to play. Such attacks have been mounted notably by Hofweber (2009), Daly (2012), and J. Wilson (forthcoming) (see Chapter 1, section 2.2). In response (as well as pre-emptively) grounding has been defended by appeal to intuitive examples of grounding, by highlighting the usefulness of grounding, and by citing connections between grounding and other, better understood metaphysical notions. Other attempts to fix the content of the term ‘grounding’ specify logical and structural features of ground (e.g. Schaffer, 2009; Fine, 2012a; 2012b; Correia and Schneider, 2012; Rosen, 2010), and/or connect the notion of grounding to that of explanation (e.g. Kim, 2004; Fine, 1995a; 2001; 2012a; Schneider, 2006; Schaffer, 2009; 2010a; deRosset, 2010; 2013; Audi, 2012a; 2012b).

In the next two chapters I present my own challenge to the orthodox conception of grounding. The first of these (which I pursue in this chapter) is concerned with the logical and structural properties of grounding. The second (which I discuss in Chapter 4) deals with the claim that we can understand the

notion of ground by appeal to that of explanation. In the final chapter, I will argue that my own antirealist approach to grounding is consistent with the upshot of both of these challenges.

My target in this chapter is the common assumption that grounding is an asymmetric relation. My strategy for calling this assumption into question is first to highlight some situations that I think are best described by taking grounding instead to be a nonsymmetric relation, and second to develop an approach to metaphysical structure that rejects the asymmetry assumption, and to argue that this approach has some advantages over the more standard conceptions of metaphysical structure. In my view, this new position is a worthwhile addition to the taxonomy of positions in the debate over metaphysical structure as described in terms of grounding.

As I said in the first chapter I will stay neutral, so far as possible, about the relata of the grounding relation and about whether we ought to formalise grounding as an operator or a predicate.²² I will assume that every entity in the domain(s) under discussion (be it facts, objects, states of affairs etc., or any combination) enters into grounding relations. Note that my use of expressions such as ‘ x grounds y ’ is not intended to exclude the possibility that there are further entities that also ground y , and so the reader might instead read ‘partially grounds’ if that locution is preferred. In this chapter I use ‘ x grounds y ’, and ‘ y (ontologically) depends on x ’ synonymously.

Recall that grounding relations are usually taken to be transitive, irreflexive and asymmetric (e.g. Cameron, 2008; Rosen, 2010; Schaffer, 2010a; Clark and Liggins, 2012: 817). This chapter develops a particular theory of metaphysical structure generated by dropping the asymmetry constraint on grounding. I call this theory ‘metaphysical interdependence’ (or ‘interdependence’ for short). Interdependence is an alternative to both the traditional foundationalist view, and the infinitist conception of metaphysical structure already discussed in the literature (and described below). Whilst the possibility of symmetrical grounding has been raised by some (e.g. Fine 2012a; 2012b on ‘weak ground’), with the exception of Bliss (2011) (who develops a detailed account of a cluster of positions she calls

²² For ease of formulation I will use the locution ‘ x grounds y ’, though the reader may substitute this for ‘ y because x ’ if s/he so desires.

‘metaphysical coherentism’) there is little defence of the possibility of symmetrical grounding, or serious discussion of resultant theories of metaphysical structure.

1. Foundationalism, infinitism, interdependence

In most recent literature concerned with fundamentality as characterised by grounding relations, two theories of metaphysical structure dominate. These two theories can be distinguished by what they have to say about well-foundedness:

Well-foundedness: for all x , x is either grounded by some fundamental entity or entities, or is itself a fundamental entity.

To facilitate discussion of the position I wish to introduce here there is a second thesis by which we might further distinguish theories of metaphysical structure:

Asymmetry: for all x and all y , if x grounds y , and $x \neq y$, then y does not also ground x .²³

Recall that following Fine (2012a: 50), we can distinguish between full and partial grounds. x is a full (or complete) ground for y if x , on its own, provides the entire metaphysical explanation for y . No other entity than x is required for y to be fully grounded. x is a partial ground for y if x , by itself or with some other entities, is a ground for y (see Chapter 1, section 2.1). In this chapter, ‘grounds’ is to be read as ‘partially grounds’, and so asymmetry dictates that no entity y can be a ground for x (i.e. can be a member of the set of entities that jointly comprise a complete ground for x) if it is the case that x appears amongst the grounds for y . Asymmetry therefore guarantees that grounding relations only run in one direction, and thus allows that entities can be related in linear chains of grounding. Asymmetry precludes any non-linear arrangement of entities in grounding chains. Once we admit of a grounding

²³ That x and y are not the same entity is built into my definition of asymmetry here so that my characterisation of the positions in the debate is acceptable both to those holding that grounding is irreflexive, and for those (e.g. Lowe, 2010 and Jenkins, 2011) who do not wish to rule out the possibility of reflexivity. Others sometimes use the term ‘antisymmetric’ to pick out a relation such that whenever x and y are distinct entities, whenever x is R-related to y , y cannot be R-related to x . Nothing turns on my choice of terminology here.

relation that applies to every entity in the relevant domain (as assumed here), the possible combinations of attitudes towards these two theses give rise to four main positions.²⁴

Acceptance of both well-foundedness and asymmetry leads to what can be considered the orthodox position in the debate over metaphysical structure: metaphysical foundationalism. Foundationalists hold that grounding chains terminate in one or more fundamental entities,²⁵ and as we move up the chain from the fundamental we encounter entities that can be considered increasingly derivative. Since grounding is transitive, the fundamental entities collectively provide the ‘ultimate ground’ for reality. The collection of fundamental entities can thus be thought of as comprising reality’s ‘fundamental level’.

The metaphysical foundationalist position is analogous to foundationalism in epistemology. Foundationalist theories of epistemic justification hold that all justified beliefs are either properly basic beliefs, or are justified inferentially through linear chains of belief that terminate in properly basic beliefs. Properly basic beliefs are those for which no further justification is required (canonical examples include *a priori* beliefs and appearance beliefs, see e.g. Fumerton, 2006). The debate about the structure of justification can help shed light on the less-well-understood debate about metaphysical structure, and so the analogy is one I will return to throughout this chapter.

Metaphysical foundationalists fall into two camps: pluralist and monist. While pluralists hold that there are multiple fundamental entities, monists posit only one. For example, a familiar pluralist foundationalist conception holds that everything that exists depends on the plurality of fundamental mereological atoms – particles that have no proper parts. The recent champion of monist foundationalism is Jonathan Schaffer (e.g. 2010a; 2013) who holds that the cosmos is the single fundamental entity, on which everything else depends.

The second position, where asymmetry is accepted but well-foundedness is rejected, is called metaphysical infinitism (e.g. Morganti, 2009). Accepting asymmetry generates chains of grounding where each entity appears only once in the chain, but rather than bottoming out in some fundamental entity,

²⁴ For ease of discussion, we can assume at this point that proponents of each view of grounding take their position to hold with necessity. I later discuss some complications introduced by contingency.

²⁵ Where what constitutes a foundational entity will depend on one’s own view of fundamentality. For the purposes of this chapter, I will assume (with the majority) that an entity is fundamental when nothing grounds it.

these chains continue infinitely. Defenders of such a view might, for example, hold that reality is infinitely divisible (the world is ‘gunky’) and that entities at each ‘level’ of reality are grounded by entities in the level below. Suppose, for example, that macroscopic objects are composed of and ontologically depend on molecules, which in turn depend on atoms, themselves dependent on sub-atomic particles, which depend on smaller sub-atomic particles, and so on ad infinitum.

Note that the above picture is not a straightforward consequence of rejecting well-foundedness and accepting asymmetry. All that is required for well-foundedness to be false is for there to be at least one entity that is neither fundamental, nor dependent on a fundamental entity. We might call the position generated by rejecting well-foundedness and accepting asymmetry ‘weak infinitism’, to be contrasted with the ‘strong infinitism’ that is the consequence of endorsing asymmetry along with the converse of well-foundedness:

Non-well-foundedness: for all x , x neither is grounded by some fundamental entity or entities, nor is itself a fundamental entity.

For simplicity, my discussion here concerns strong infinitism.

Again, there is a helpful analogy with infinitism in case of epistemic justification. Infinitists about justification hold with the foundationalist that justification is inferred through linear chains of reasoning, but maintain that any termination of the chain is arbitrary. There is no last reason in the series, and no belief can appear twice in the same chain of reasons. Beliefs are justified when they are part of a non-terminating, non-circular series of reasons (see e.g. Klein, 2005).

The third available position is metaphysical interdependence, which rejects both well-foundedness and asymmetry. According to this view, some entities are such that they are not grounded in anything fundamental, and some entities are such that there is mutual ontological dependence between them. Imagine, for example, a world containing only a circle and a pair of its semi-circles (see figure 1 below).




Figure 1

It might be that at that world, the circle is grounded in the semicircles that compose it (a ‘metaphysical explanation’ of the circle cites the semicircles; once we have the semicircles, the circle is ‘no addition to being’) but the semicircles are also grounded in the circle – their existence derives from the existence of the circle. In this case, there is a symmetrical grounding relation between the circle and the semicircles. Since everything at this circle-world is grounded in something else, nothing is ontologically basic; nothing is fundamental. (See the following section for some more detailed examples of symmetrical grounding.)

When there are more objects at a world, the rejection of well-foundedness and asymmetry allows for the possibility of complex grounding networks and grounding loops. Once again, there is a familiar epistemic analogue of this view. Coherentists about justification take a belief to be justified when it is part of a coherent system of beliefs. The coherentist belief system is a network of mutually supporting beliefs where a given belief can appear in its own reason ancestry (and justification is an emergent feature of such networks). Metaphysical interdependence, like coherentism about justification, posits networks of entities where there is no prohibition on entities partially depending on themselves. This is demonstrated in the above case of symmetrical grounding involving circles and semi-circles.

But again, just as coherentists posit complex networks of beliefs where a belief usually appears far away from itself in its own reason ancestry, metaphysical interdependence too recognises that symmetrical grounding may not be immediate. Suppose we have four entities: A, B, C and D, related in such a way that A grounds B, B grounds C, C grounds D and D, in turn, grounds A. Since grounding is transitive, it will also be the case that A grounds, albeit mediately and partially, C, D, and itself. It follows then that if

grounding is nonsymmetric and transitive, it is also nonreflexive (because it must be possible for something to ground itself). Figure 2 below illustrates this. (The direction of the arrow runs from grounding entity to grounded entity. The purple lines represent the immediate grounding relations, and the red lines represent the mediate grounding relations that are the consequence of transitivity.)




Figure 2

As with infinitism, we can distinguish weak and strong versions of interdependence. On the weak conception (as outlined above) there might be some fundamental entities and some parts of linear chains of dependence. A strong version of interdependence endorses non-well-foundedness in conjunction with the converse of asymmetry:

Symmetry: for all x and all y , if x grounds y , then y grounds x .

The grounding between x and y might well be partial and indirect, but that symmetry entails that for all x and all y , if the metaphysical explanation for y includes x , then the metaphysical explanation for x will include y . This claim is weaker than it perhaps first sounds. Somewhere in the metaphysical explanation of y we will find x , but it might be that the metaphysical explanation includes all the entities at the world where x and y exist (compare the occurrence of a reason Q in a coherentist's belief network that is only cited in justification of some belief P (which is itself partly justifies Q) if every reason in the network is cited in justification of P).

(Note that it is consistent with the strong version of interdependence that all entities are grounded in mutually dependent pairs, and no grounding relations obtain between each pair and any other entity. This strikes me as an unlikely way for the world to be, and it is unproblematic for metaphysical interdependence in the same way as it is unproblematic for the foundationalist that her theory would allow for similarly isolated grounding pairs so long as one member of the pair is fundamental. There are good reasons in both cases to think that the world does not conform to those models.)

For completeness, we should note that a fourth position accepting well-foundedness and rejecting asymmetry is *prima facie* possible. A weak version of this view has it that structures of grounding are always well-founded, but that there might be some interdependent loops in the well-founded chains. This is very similar to the weak version of interdependence, though less permissive because it requires rather than merely permits that networks of grounding relations are well-founded. A strong version of this fourth position accepts symmetry and well-foundedness, and thus entails that there can be no instances of grounding (whenever two or more entities enter into grounding relations, the grounding between them must be symmetrical, but this has the consequence that there are non-well-founded entities). Since this position precludes any instances of grounding, I do not discuss it further.

2. Examples of nonsymmetric grounding

The claim that grounding relations are nonsymmetric may seem *prima facie* implausible. This section highlights extant examples that I think are best interpreted as cases of symmetric grounding, in order to help dispel that impression. Those who do not buy into the assumptions made in this section may still find the examples somewhat plausible and thereby soften their attitudes towards nonsymmetric grounding. The examples considered involve (i) propositions; (ii) essences; (iii) quantities; (iv) structural realism; and (v) holism.

Propositions: A first example concerns propositions that are such that they depend mutually on one another for their truth. Assume that propositions are grounded in their constituents, and are either true or false. The proposition <the Earth is flat> is grounded in the Earth, and the property of flatness. Similarly, the proposition (call it P_0) <the proposition that the Earth is flat is true> is grounded in the

proposition that the Earth is flat, and the property of being true (assuming that there is such a property). Now take the proposition $\langle P_2 \text{ is true} \rangle$ and call it P_1 . P_2 is the proposition $\langle P_1 \text{ is true} \rangle$. P_1 is grounded in P_2 , and in the property of being true. P_2 is grounded in P_1 (and the property of being true) and so these two propositions exhibit a symmetrical grounding relation between them. (If P_1 is false, its falsity will be grounded in P_2 , which is itself grounded in P_1 .)

Essence: Kit Fine (1995b) offers an ontology of the essences of fictional characters where he considers the possibility of the essences of two such characters mutually depending on one another. It might, for example, be of the essence of Holmes to be admired by Watson, and of the essence of Watson to admire Holmes (Fine, 1995b: 65). In this case we can say that the essence of Holmes is (partially) grounded by the essence of Watson, and vice versa. Whilst it is possible to avoid this mutual grounding (we can say that it is essential to both Holmes and Watson as a pair that one admire the other) Fine (1995b: 66) suggests that the mutual dependency relation interpretation might be preferable. Rather than supposing that some objects enter into particular relations essentially (in this case, the admiration relation), we can explain mutual dependency solely in terms of the essences of the relevant objects themselves.

Quantities: A third example also comes from Kit Fine (2001: 11).²⁶ The mass, density, and volume of a portion of a homogenous fluid are interrelated in such a way that any the value of any two of the three parameters might be said to ground the value of the other. There appears to be no principled reason for taking any one of the three parameters as derivative of the other two. This constitutes a violation of either the transitivity of grounding, or of the asymmetry and irreflexivity of grounding. The figure below illustrates the (immediate) grounding relations between these parameters (the direction of the arrow is from grounding entity to grounded entity).

²⁶ Fine is not arguing for interdependent grounding with this example. Instead he is arguing against the idea that reduction is a matter of supervenience on the grounds that we cannot say, without circularity, that each parameter is reducible to the value of others.

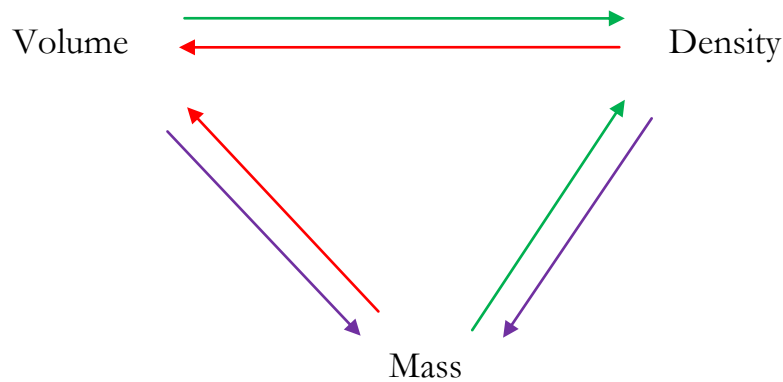


Figure 3

If grounding is transitive, then as figure 3 above demonstrates, each of the parameters (partially) grounds both the other two, and itself. If grounding is interdependent, we can embrace the apparent circularity and hold that each of the three parameters is indeed grounded in the other two. This solution is the best fit with our understanding of the relations between these quantities, which seems to dictate that in a given circumstance, we might take any two parameters to ground the third.

Structural Realism: A fourth case of interdependent grounding is found in a recent proposal in the philosophy of science. Structural realism is a now-popular form of scientific realism which holds that rather than asserting that the nature of things is correctly described by our best scientific theories, the realist emphasis should be on the structural content of those scientific theories. Scientific realism was introduced into contemporary philosophy of science by John Worrall (1989), but was arguably anticipated by Russell (1927) (see Demopoulos and Friedman, 1985). The motivation for the position is that it is structural content that is retained across theory change (Ladyman and Ross, 2007: 93). Structural realists of a metaphysical variety take structure and relations not to be merely derivative of the entities or structural nodes they relate, but to be more ontologically fundamental than has traditionally been assumed by scientific realists (Ladyman, 2009). Esfeld and Lam (2006) develop a version of metaphysical structural realism that they term ‘moderate structural realism’ and which holds that relations require relata, but that it is not the case that the relata necessarily have intrinsic properties over and above the

relations they bear to one another (Esfeld and Lam, 2006: 32). In other words, there are objects, but those objects are wholly characterised by the relations in which they stand.

It follows from this characterisation of moderate structural realism that there is a mutual dependence between relations and relata – the objects are characterised by the relations that relate them, and the relations themselves require the existence of the objects that stand in the relations. There is therefore:

‘...a mutual ontological as well as conceptual dependence between objects and structure (relations): objects can neither exist nor be conceived without relations in which they stand, and relations can neither exist in the physical world nor be conceived as the structure of the physical world without objects that stand in the relations.’ (Esfeld and Lam, 2006: 32.)

If moderate structural realism is plausible, so are symmetric instances of grounding.

Holism: More of an example of metaphysical interdependence in action than just of symmetrical grounding, the doctrine of holism states that systems ought to be evaluated as wholes rather than as collections of their parts. Whilst in some areas of philosophy holism remains unpopular, in others there are well developed holist theories which are taken seriously (e.g. confirmational holism; holism about meaning). Esfeld (1998) argues that a holistic system is characterised by a mutual ontological dependence among its parts, and that for that ‘we have to conceive generic ontological dependence as a symmetric relation’ (1998: 369). If ontological dependence is to be expressed in terms of grounding, holist doctrines thus require a theory of grounding that allows grounding relations to be symmetric.

Any or all of the assumptions required to generate the examples above might be false, but none of them are wholly implausible. Since each of them either requires or results in nonsymmetric instances of grounding, I take the examples to demonstrate that nonsymmetric grounding is also not wholly implausible. In the next section, I consider some positive arguments for metaphysical interdependence.

3. Arguments for metaphysical interdependence

This section provides four arguments motivating acceptance of an interdependent theory of metaphysical structure: the first based on intuitions concerning mereology; the second on the permissiveness of different theories of metaphysical structure; the third on considerations about theoretical virtue; and the fourth on a puzzle about how grounding is itself grounded.

3.1. Part-whole grounding

The argument of this sub-section is that metaphysical interdependence is the only theory of grounding able to reconcile competing intuitions about the dependency relations between parts and wholes. It is common in the literature to hold that the relation between parts and wholes should be considered a relation of grounding (e.g. Correia, 2008; Schaffer, 2003; 2010a). Recall that priority monists are foundationalists holding that there is exactly one fundamental entity: the entire cosmos (Schaffer, 2010a). Arguments for priority monism uncover compelling intuitions about the priority of wholes over parts. It is natural to suppose, for example, that semi-circles (such as those in figure 1) are grounded in their circles, or that an organism might be prior to its organs (Schaffer, 2010a: 47). The existence of the semi-circle seems to derive from the existence of the circle; without the circle there would be no semi-circles. It seems as though it is in the nature of the organs of an organism to play the role they play as parts of that organism – their nature seems to derive from the nature of the organism.

Arguments from emergence (see Schaffer, 2010a: 50-57) suggest that we cannot tell a complete story in terms of parts alone. Whilst the metaphysics of entanglement remains controversial, a dominant view of quantum-mechanically entangled systems is that ‘properties of the whole contain all that can be said about the local properties of the parts, and *only* these properties of the whole contain all that can be said about the local properties of the parts’ (Esfeld, 2001: 252, quoted in Schaffer, 2010a: 52). If there are any genuine cases of emergence, they are cases where it is not possible to reduce the properties of the whole to properties of its constituent parts.

Nevertheless, the intuition that wholes depend on their parts is hard to shake off. Whilst it might be the case that there are features of a system that cannot be explained purely in terms of the parts of the

system and the relations between them, it is counterintuitive to think that the system does not depend on its parts at all. After all, the parts are what the system is composed of – without the parts, there would be no whole.

Whilst the existence of the semi-circles seems to derive from that of the circle, the circle is nevertheless made up of the semi-circles. The truth of the proposition <the semi-circles exist> seems to consist at least partly in the truth of the proposition <the circle exists>. But the truth of the proposition <the circle exists> seems to consist (at least partly) in the truth of the proposition <the semi-circles exist>. Similarly, it seems as though whilst the organs of an organism derive their nature from the organism, the organism requires its organs in order to exist and to be that organism. Metaphysical interdependence offers a way to reconcile these competing intuitions. If the grounding relations between parts and wholes are symmetrical, then we are able to consistently maintain both that wholes depend on their parts, and that parts depend on their wholes.

3.2. Gunk and junk

The second argument for metaphysical interdependence again draws on the claim that that the relation between parts and wholes is one of grounding, and concludes that only metaphysical interdependence is capable of reconciling the metaphysical possibilities of both ‘gunky’ and ‘junky’ worlds.

Gunky worlds are those at which everything has proper parts. Junky worlds are those at which everything is a proper part of something. Both junk and gunk appear to be metaphysically possible (see discussions in, for example, Sider 1993; Hudson 2007; Varzi 2009; Schaffer 2003, 2010a; Bohn 2009a; 2009b) but foundationalists will be unable to account for both possibilities. Other things being equal, a theory about grounding that is compatible with both gunk and junk will be preferable, but foundationalists are unable to provide such a theory. This criticism is strongest against the foundationalist who takes foundationalism to be necessary if true. (It is open to the foundationalist to claim that foundationalism might be merely contingently true (as in Cameron, 2008), and thus that the mere metaphysical possibilities of gunk and junk are not troubling. In that case the onus is on the defender of foundationalism to explain why we should expect these metaphysical theses to hold

contingently. In any case, my argument here can be restricted to the actual world, where there are not yet decisive reasons to think that the world is not gunky, junky, or both ('hunky').)

Pluralist foundationalists holding that the multiple fundamental entities are whatever is the smallest unit of reality (e.g. mereological atoms) have a problem accounting for gunky worlds. In gunky worlds everything is divisible, and thus there will be no smallest unit of reality, and so no fundamental entities to ground reality. The pluralist foundationalist in this scenario holds that each entity depends on its parts, and since there is no entity that does not have any parts, there are no non-dependent entities; everything is grounded in something. Since the foundationalist holds that everything is either fundamental (i.e. ungrounded) or grounded in the fundamental, gunky worlds are incompatible with their theory.

It is at least *prima facie* possible to be a pluralist foundationalist and hold that the fundamental entities are not those with no parts, but might be the medium-sized dry goods with which we are familiar from our everyday interactions with the world. A foundationalist of this stripe holds that all of the rest of reality depends on those fundamental macroscopic entities for its existence and nature, and this foundationalist therefore does not have a problem with gunky worlds (everything is grounded in the medium-sized dry goods). Reality might be infinitely divisible, but grounding chains nevertheless terminate, radiating inwards towards the familiar macroscopic objects. This kind of view might be defensible, but it is counterintuitive and would require independent motivation. Since I am not currently aware of it having any defenders, I set it aside here.

Schaffer (2010a: 61-64) argues that whilst the pluralist cannot account for gunky worlds (which he takes to be both metaphysically possible and consistent with current science) the foundationalist monist can. The monist holds that the entire cosmos is the single basic object in which all of the parts of the cosmos are grounded. Consequently, even if everything has proper parts, all of these parts are grounded in the ultimate whole. But just as the pluralist foundationalist cannot account for gunky worlds, the monist foundationalist will be unable to account for the possibility of junky worlds, because at such worlds there will be no ultimate whole in which to ground reality. Everything is a proper part of some larger object.

In a nutshell, the metaphysical foundationalist will always be unable to simultaneously account for both gunk and junk. This leaves her two options. She must (1) adopt either monism or pluralism and then

explain why either gunk or junk respectively is metaphysically impossible (this is the strategy adopted in Schaffer, 2010a); or (2) argue that whilst both gunk and junk are contingent possibilities, all the worlds at which pluralist foundationalism is true are non-gunky, and all worlds where monist foundationalism is true are non-junky (and no foundationalist worlds are hunky). In response to the first option, I have suggested that it is a mistake to deny the metaphysical possibility of either gunk or of junk. In response to the second, we can note first that most take their view of metaphysical structure to be non-contingent, and second that that we could restrict our argument to the actual world. Our best current science does not currently militate against gunky worlds, and the scientific debate over the finitude of the universe and philosophical debate over composition are both live. The foundationalist who adopts the second strategy is hostage to empirical fortune over the outcome of those debates.

In contrast, metaphysical interdependence has no problem with worlds that are either gunky, junky, or hunky. Interdependence necessitates neither a smallest particle nor a largest sum, as grounding relations do not form a linear ordering. Where the relevant debates are unresolved, it seems prudent to adopt a system of grounding that will be consistent with any development in those debates, both concerning the metaphysical possibilities of gunk, junk and hunk, and concerning the structure of the actual world.

Metaphysical interdependence is better placed than foundationalism to deal with all of the metaphysically possible ways in which reality might be structured. It is more liberal than infinitism, which apparently *requires* that the world be gunky or junky in order to populate its infinitely extending grounding chains. Interdependence can cope with finite or infinite levels of reality, or dispense with them altogether. It can thus remain neutral on these issues, and acceptance or rejection of the view will not turn on any future discovery, scientific or otherwise, about whether reality ‘bottoms out’ or composition is unrestricted. Interdependence is non-arbitrary in the sense that it does not have to make any stipulations in these areas.

3.3. Theoretical virtue

A third argument for metaphysical interdependence is that, with respect to the three virtues considered below (unity; parsimony; and commitment to brute facts) metaphysical interdependence is more

theoretically virtuous than either infinitism or foundationalism. What does or doesn't count as theoretically virtuous, questions about the argumentative weight we should attach to theoretical virtue, and the issues involved in evaluating rival claims to virtue are at best murky, and so arguments based on theoretical virtue are likely to be suggestive rather than conclusive. Nevertheless, our discussion here concerns some of the less contentious examples of theoretical virtue, and that interdependence performs better than rival theories gives us at least some reason to believe that metaphysical interdependence is true.

First Virtue – unity: It is more theoretically virtuous, and thus preferable, if when explaining some phenomena, we can give the same explanation of each phenomenon rather than separate explanations of each (Cameron, 2008: 12). Imagine getting up one morning and discovering that power lines are down, trees have fallen across the road and cars are overturned. There are two candidate explanations: (i) there was a storm in the night; or (ii) the street was visited overnight by malicious power-line workers, reckless tractor drivers and gangs of car-flipping children. (i) provides a better explanation of the destruction observed than (ii) does because it is a single, unified explanation for all the observed phenomena.

Cameron (2008: 12) appeals to explanatory unity in order to argue that foundationalism is more theoretically virtuous than infinitism, and thus is the correct account of metaphysical structure. In the background here is the common assumption that when y grounds x , y provides a kind of metaphysical explanation for x . Cameron's argument is that if infinitism is true, whilst each individual component of reality is grounded, there is 'no explanation of everything that needs explaining' (2008: 12). The metaphysical explanation (or ground) for each entity is the entity below it in the grounding chain, but there is no unified explanation of every dependent entity. Cameron then claims that if, on the other hand, there is some set of ontologically fundamental entities, then there is a unified ground, and hence a unified explanation, for everything that exists.

It is, however, open to the infinitist to reject Cameron's characterisation of the matter. First, she is likely to deny that it is in any sense better to give a single metaphysical explanation for every entity, rather than some explanation for each entity. What matters is that everything is grounded, and not that everything is grounded in the same entity or set of entities. But even granting Cameron's assumption, the

infinetist can argue that there *is*, on her view, a unified explanation of everything that needs explaining – there is the entire infinitely extending chain of grounding; everything explains everything. In order for Cameron’s argument to work, we require a characterisation of unity that prohibits this kind of infinitist response.

Intuitively, when the infinitist appeals to the entire grounding chain to give a unified metaphysical explanation of any given entity in the chain, many of the entities she appeals to in her explanation are *superfluous*; an adequate explanation could have been given which didn’t mention those entities. Getting clear on definitions of superflousness and unity will allow us to assess how well metaphysical interdependence compares with foundationalism and infinitism with respect to theoretical unity. We can define superflousness in this context as follows:

***y* is superfluous to grounding the *x*s** =_{df} there are some *x*s and some *z*s such that the *z*s ground the *x*s, *y* is one of the *z*s, and the plurality of the *z*s without *y* provides a complete ground for all of the *x*s.

An entity *y* is thus superfluous when *y* could be eliminated from the grounding explanation (i.e. the specification of the full ground for every entity in the system), and all of the *x*s still be explained. In other words, *y* is superfluous when a complete ground for the *x*s need not include *y*. For Cameron’s argument for foundationalism to go through, he must have in mind that unified metaphysical explanations do not include superfluous entities (otherwise he would not have grounds for thinking foundationalism more virtuous on grounds of unity than infinitism).

Following Quine (1948), we can take the ontological commitments of a theory to be whatever must exist in order for the theory to be true. We can now define unity, in the way that Cameron must understand it, as follows (where *T* is a theory of metaphysical structure):

***T* is unified** =_{df} for all *y*s that *T* ontologically commits to, there are some *x*s that ground the *y*s, and none of the *x*s is superfluous.

With these definitions in place, we see that foundationalism is unified because there is a set of fundamental entities that collectively ground everything else that exists, and none of these entities is

superfluous (none of these fundamental entities can be left out of a complete explanation of reality). In finitism is not unified because although for the finitist the entire chain of entities forms a set of entities that ground everything, any entity we care to consider will turn out to be superfluous to a grounding explanation. For any given entity, its grounding work can be done instead by the entity below it in the infinite grounding chain. If x grounds y and y grounds z , then z would still be grounded without y , because it would be grounded by x .

Whilst finitism comes out as disunified on this account of unity, metaphysical interdependence *is* unified in accordance with the above definition. Entities involved in webs of dependence are such that when any entity grounds another, it also grounds every other entity in the web, and itself.²⁷ Everything thus explains everything, but unlike in the finitist case, no entity is superfluous. No entity can be left out of a complete grounding explanation, because each entity is needed in order to provide part of the complete ground for all of the other entities in the web. If x grounds y , y grounds z , and z grounds x , then when y is removed from the system we can no longer give a complete ground for x or z . This is easiest to see by means of an example. Let x be the proposition $\langle y \text{ is true} \rangle$; y be the proposition $\langle x \text{ is true} \rangle$ and z be the proposition $\langle z \text{ is true} \rangle$. The truth of each of these propositions is grounded in the other two, but removing y leaves the truth of the other two ungrounded. Like foundationalism, metaphysical interdependence is unified in accordance with the sort of definition of unity given above.

Second virtue – parsimony: In this context, a more parsimonious theory postulates fewer entities, or kinds of entities, than its rivals. Other things being equal, theories that commit us to smaller ontologies rationally ought to be preferred (Baker, 2013). There are ways in which each of foundationalists, finitists and defenders of interdependence might make claims to parsimony, but I argue here that defenders of foundationalism and interdependence have the stronger claims.

Since finitism is committed to infinitely extending grounding chains, it is committed to the existence of an infinite number of whatever entities are related by the grounding relation. In the case of actual concrete objects, for example, finitists commit themselves at the very least to the existence of infinitely

²⁷ Recall that for ease of evaluation, our discussion here involves strong interdependence. The argument can be straightforwardly modified to apply to weak interdependence (and also weak finitism).

many proper parts of such objects in order to populate the infinite chains of dependence. Neither interdependence nor foundationalism is committed to the existence of infinite numbers of any type of entity, and each can thus claim to be more parsimonious than infinitism.

A foundationalist can choose to populate her ontology with only those entities which she takes to be ontologically fundamental (see Chapter 1, section 3.2). In the case of actual concrete objects, for example, this ensures that whilst the world is populated by the everyday entities of common sense and science (chairs, tables, people and the like), the foundationalist's ontology remains sparse.²⁸ Where defenders of infinitism and interdependence must take their ontology to be comprised of all of the entities at a world, the foundationalist need commit herself only to the fraction of those that are fundamental. She can therefore claim that her theory is more parsimonious than those of her rivals.

A defender of interdependence has grounds to challenge the foundationalist's claim to parsimony. A quick (and dirty) response to the parsimony-claiming foundationalist is as follows: Only fundamental entities carry ontological commitment; interdependence holds that there are no fundamental entities (assuming all entities enter into grounding relations); ergo defenders of interdependence have an empty ontology, and it doesn't get more parsimonious than that!

But there is a further argument against the foundationalist's claim to parsimony that appeals to the distinction between foundational and non-foundational entities. We should be concerned not only with quantitative, but also with qualitative parsimony. Where the former counts instances of a given entity, the latter counts types of entity. There is a sense in which the foundationalist countenances an extra *type* of entity, i.e. those that are fundamental, and thus carries an extra qualitative commitment. If we keep score in this way, friends of interdependence and infinitism do better, as they have no room for a category of *fundamentalia* in their ontology. (Clearly, more needs to be said if this is to be taken to be a virtue of the theory in any further sense. The argument here concerns only judgements about parsimony.)

Of course, the foundationalist can quickly respond that their ontology too includes only one type of entity. Only entities of the fundamental type are ontologically committing. The defender of

²⁸ This is at least true of the majority of foundationalists, though some (e.g. Schaffer, 2010a; Audi, 2012) deny that non-fundamental entities are in any sense 'less real' than the *fundamentalia*.

interdependence can legitimately level a charge of question-begging against the foundationalist here (that only fundamental entities are ontologically committing is foundationalist dogma), but he cannot insist that qualitative parsimony be measured on interdependent terms lest the same charge be brought against him. We have a kind of stalemate between the defender of foundationalism and the defender of interdependence, at least pending an independent resolution of the metametaphysical debate concerning ontological commitment.

Third virtue – brute facts: Brute facts are those for which no explanation can be given. The more brute facts a theory is committed to, the more things it leaves unexplained, and thus the less virtuous the theory is. Foundationalism is committed to brute facts about fundamental entities (by definition, such entities cannot be explained by anything further). Infitism and interdependence require no such commitment to brute facts; facts about each entity can be explained by appeal to facts about further entities.

The foundationalist might respond that facts about fundamental entities are not unexplained, but rather self-explanatory. (This is analogous to foundationalism about epistemic justification, where foundationalists hold that the properly basic beliefs from which justification for all other beliefs is inferred are self-justifying.) To say that facts about fundamental entities are self-explanatory is to come very close to say that fundamental entities are self-grounded. (Indeed, if the relata of the grounding relation are facts, it is to say that the fundamental facts ground themselves). This is tantamount to a rejection of the claim, characteristic of foundationalism (cf. Schaffer, 2010a; Rosen, 2010; Audi, 2012a; 2012b), that grounding is an irreflexive relation. In fact, the only theory of metaphysical structure that embraces nonreflexivity in grounding is metaphysical interdependence. By transitive closure, entities related in an interdependent way end up being partially self-dependent (see figure 2 above), and thus reflexivity is a feature of interdependence. To claim that facts about fundamental entities ground themselves will not help the foundationalist, and so the charge of a non-virtuous commitment to brute facts remains.

Foundationalism and interdependence share the virtues of both parsimony and explanatory unity, though they realise them in different ways. Interdependence and infinitism share the virtue of carrying

no commitment to brute facts. Since interdependence exhibits both virtues, it is the most virtuous of the theories under consideration. This is motivation to accept interdependence over other theories of grounding.

3.4. Grounding grounding

The final argument for metaphysical interdependence I will discuss here concerns the ability of the defender of interdependence to respond to a question posed in Bennett (2011) (also addressed in Sider, 2011 and deRosset, 2013) about whether grounding is itself grounded. For the foundationalist, this question becomes a dilemma about whether or not the grounding relation is fundamental. Bennett argues that if grounding is not grounded (i.e. is fundamental) then two compelling principles about grounding are violated.

The first is the purity principle argued for in Sider (2011: 115-6), which entails that connections between the fundamental and the non-fundamental cannot themselves be fundamental. The second is Schaffer's (2010a: 40) principle that the fundamental elements of the world should be open to free modal recombination – since ungrounded entities do not depend on anything else, they should be modally unconstrained. Bennett argues that if grounding is fundamental, then there is a possible world w just like this one in terms of the distribution of all the rest of the fundamental entities, except that nothing grounds anything else. Any actually grounded entity must, in w , either fail to exist or be fundamental, and neither option is plausible (cf. Bennett 2011: 27).

On the second horn of the dilemma, grounding is grounded. The problem with this solution is easiest to see when put in terms of grounding facts (though see Bennett's paper for a discussion in terms of relations). If x grounds y , and grounding is grounded, then we can ask what grounds the fact that x grounds y . If this is some further fact (call it A), then we can ask what grounds A . Regress threatens, but if foundationalism is true then the regress is intolerable. Bennett's proposed solution is to grasp this second horn of the dilemma, but to argue that the regress generated is benign. For each iteration of the question about what grounds the grounding fact, Bennett (2011: 34) contends that the answer is the first relatum of the grounding claim. It is x that grounds the fact that x grounds y , x that grounds the fact that

grounds the fact that x grounds y , and so on. There is a regress of grounding facts, but the regress is itself grounded in a way that is compatible with foundationalism.²⁹

I find the bootstrapping in Bennett's proposal unsatisfying. When we ask for the ground of the grounding fact, we want to know in virtue of does grounding relation between x and y obtains. To learn that it is x that grounds that fact, and furthermore that it is x that grounds the fact x grounds that fact leaves us feeling as though we have not been given a proper explanation. Compare a child who asks why she has to go to bed, and is told 'because I said so'. If she asks 'but why do I have to go to bed because you said so?' and again receives the answer 'because I said so' -- we can sympathise with her indignation. I feel a similar indignation in the face of Bennett's answers to questions about grounding grounding.

Indignation is of course no substitute for argument, but I think its root is in the feeling that the explanation offered is somehow incomplete. This is best seen by way of an example. Disjunctions are said to be grounded in their true disjuncts, so, when P is true, the disjunction $P \vee Q$ is grounded in P . When we ask what grounds the grounding fact, Bennett replies again ' P '. But this doesn't seem right. It is not just P that grounds the grounding fact, but P along with a further fact about disjunction: that the truth of a single disjunct is sufficient for the truth of the disjunction. In citing just P we give an incomplete explanation.

DeRossett (2013: 21) responds to this sort of objection by claiming that these further facts should not be considered part of the explanans. It might be that an adequate explanation must make its explanandum intelligible to an audience, and that in order to do so more features of the explanatory story than just the explanans must be mentioned. When we apply this to the example above, we might think that that an adequate explanation must make intelligible the relation that holds between P and $P \vee Q$, and that P by itself is not sufficient to do so. In order to make that relation intelligible we may have to cite some further facts, but deRossett's point is that these are merely ancillary material, needed only to deliver the relevant epistemic payoff to an audience. They are not themselves part of the explanans.

²⁹ A similar suggestion is offered in deRossett, 2013.

Trogon (2013b) discusses a reply to DeRosset's response, which he attributes to Shamik Dasgupta. Dasgupta accepts that facts about disjunction are not part of the explanans when we metaphysically explain $P \vee Q$ in terms of P , but claims that it does not follow that they are not part of the explanans when it is the fact that $P \vee Q$ is grounded in P for which we seek a metaphysical explanation. He argues for this claim by analogy: When we prove $P \vee Q$ from P there is no mention of facts about disjunction, but when we prove in the metalanguage that P implies $P \vee Q$, that disjunction introduction is a valid inference rule must be among the premises.

The foundationalist cannot accept any proposal for dealing with the dilemma which involves a regress of grounding facts, on pain of violating the well-foundedness constraint that is characteristic of foundationalism (unless the regress is grounded in a foundational fact/entity as in Bennett's proposal). Any theory of metaphysical structure that resists the well-foundedness claim does not succumb to the original dilemma. Neither metaphysical infinitists nor defenders of interdependence have a problem with the idea that grounding is itself grounded. An example helps illustrate this.

Suppose that y is a table, the x s are the molecules that compose the table, and the x s ground y . Rather than claiming that it is the x s that ground the fact that the x s ground y , a defender of interdependence can cite the x s along with other relevant aspects of the system – perhaps the spatiotemporal arrangement of the particles that accounts for why it is *those* x s (rather than some others) that ground y , and/or the law-like generalisation that composite objects are grounded in their parts. Precisely what the answer to the question of how grounding is grounded is will be depend on one's particular approaches to grounding and to composition; but note that if we adopt interdependence, there will be no problem if some reiterations of the grounding question are answered by appeal to facts or entities that have already appeared in the chain of answers to grounding questions.

4. Conclusion

Given some plausible assumptions, there are extant cases of symmetrical grounding. Metaphysical interdependence is a theory of grounding that accommodates those examples, along with competing

intuitions about dependency. It is consistent with more of the ways reality might turn out to be than its rival theories, exhibits greater theoretical virtue than its rivals, and offers a satisfying solution to the problem about grounding grounding. All of this gives us good reason to be sceptical of the orthodox view that grounding is asymmetric, and to take metaphysical interdependence seriously in the debate over metaphysical structure.

This challenge to the orthodox conception of grounding demonstrates one way in which the details of grounding are still ‘up for grabs’. What I have described as the orthodox conception is not particularly well established, and there is room for development and improvement in how the notion of grounding is to be understood. The following chapter provides a second challenge to the ‘orthodox’ view, and the final chapter brings together the lessons of this and the next to develop a new account of the nature of grounding.

Chapter 4: Metaphysical Explanation

This chapter concerns the connection between the notion of grounding, and the notion of metaphysical explanation. As should now be apparent, this is a connection which is very often made in the literature. In the first section I clarify the connection as it is generally understood, and give some examples of how defenders of grounding appeal to the connection. In the second section, I discuss three major models of explanation (the covering-law model, the causal model, and the pragmatic approach) and explain how those models might be extended to give an account of a distinctively metaphysical approach to explanation. The third section articulates a problem, in the form of a dilemma, for those who wish to appeal to a connection between grounding and metaphysical explanation. Sections 4 and 5 grapple with each horn of the dilemma, but I argue that neither can be made attractive. This sets the stage for the final chapter of the thesis, in which I respond to the dilemma by rejecting the assumption of realism about grounding that gave rise to it.

1. Grounding and explanation

Recall that grounding claims can be expressed with a number of different locutions, such as the following:

- (a) The set {Socrates} is grounded in Socrates himself
- (b) The moral properties of an action are due to its non-moral properties
- (c) The proposition <snow is white> is true because snow is white
- (d) Disjunctions are grounded in their true disjuncts
- (e) The painting is beautiful in virtue of facts about the way it is received
- (f) Legal facts depend on social facts

Grounding claims (such as (a) to (f) above) each contain an explanatory element – it is the way that the painting is received that explains why it is beautiful, the non-moral features of an action that explain what made the action morally permissible, and so on. Dreier (2004: 35) says that ground is ‘most illuminating explanation’, and the common (foundationalist) view that chains of grounding ‘bottom out’ in a fundamental entity or entities (where such entities might be propositions) is frequently expressed as the idea that those fundamental entities ‘explain’ all the rest (e.g. Jenkins, 2013).

It is the explanatory component of grounding that distinguishes ground from purely modal notions. Recall the Finean counterexamples (e.g. 1995a: 271-2). At any worlds where Socrates exists, so does his singleton set. But this necessary coexistence fails to capture the sense in which it is the existence and nature of the singleton set that depends on that of Socrates, and not vice versa. It is supposed to be Socrates’ existence and nature that *explain* the existence and nature of his singleton, and not the existence of the singleton set that explains the existence and nature of Socrates. Explanatory asymmetries obtain in the absence of modal asymmetries, and they are accounted for by appeal to grounding relations.

A different way in which connections between grounding and explanation are emphasised is by drawing attention to purported uses of grounding locutions in ordinary language. That grounding locutions like ‘because’ and ‘in virtue of’ are sometimes used in ordinary discourse (e.g. ‘it’s wrong to hit your brother because you might hurt him’) provides evidence against the suggestion that grounding and related notions are esoteric in a way that is damaging to their plausibility (as per the sorts of arguments found in Hofweber, 2009, and discussed in Chapter 1, section 2.2.2). If grounding is a semi-technical notion sometimes expressed in folk discourse then charges against its intelligibility are harder to maintain (see e.g. Rosen, 2010). The folk locutions that have been claimed to be ways of expressing grounding claims are explanatory locutions. Appealing to the explanatory component of grounding is thought to shed light on the notion because we have an intuitive grasp on the idea of what it is for one thing to explain another, and an everyday understanding of terms like ‘because’.

A third connection between explanation and grounding concerns the logical and structural properties of the two notions. Like explanation, grounding is usually supposed to be transitive; irreflexive; asymmetric;

non-monotonic; and hyperintensional.³⁰ The formal features of explanation and grounding are so similar that Daly (2012) launches an attack on the intelligibility of grounding on the basis that those features, being the same as those of explanation, don't fix its content. Note that it is the formal features of the better-understood notion of explanation that informs our understanding of those of grounding (see e.g. Rosen, 2010). This is to be expected, the argument goes, because grounding is an explanatory notion. We might therefore be tempted to take the connection between ground and explanation to be instrumental in introducing and explicating the notion of grounding to those unfamiliar with it.

One further piece of evidence that the connection between grounding and explanation is a very strong one: Fine (2001: 21-2) argues that questions of ground are settled in two ways – we can consult our intuitions about what grounds what on a case by case basis, and we can appraise systems of ground as explanatory schemes.³¹ The latter is legitimate because 'the relationship of ground is a form of explanation'. Recourse to explanatory notions is extremely useful in making judgements of ground, because we judge questions of ground in part by considering how our answers fit in to a more general pattern of explanation.

So, to say that y grounds x is to provide some kind of explanation of x in terms of y . Defenders of grounding usually insist that the kind of explanation on offer is somehow 'metaphysical'. For example, Fine introduces the idea of ontological ground as 'a distinctive kind of metaphysical explanation, in which explanans [the explanation of the explanandum] and explanandum [the phenomenon to be explained] are connected...through some constitutive form of determination' (Fine, 2012a: 37). As Fine is one of few philosophers who has something to say about the explanatory character of ground, it is worth quoting him at length. Fine says that the relation of ground is to be distinguished from other explanatory connections among truths by being the tightest such connection:

³⁰ My arguments for dropping the asymmetry constraint on grounding are not in conflict with this point. First, my argument here concerns the orthodox understanding of the two notions. Second, largely as a consequence of my sympathy towards a pragmatic approach to explanation (described in section 2.3 of this chapter) I think explanation is also nonsymmetric. Unfortunately, I am unable to discuss the logical properties of explanation in detail here.

³¹ I am a little unclear what Fine means by an explanatory scheme in this context, but his suggestion seems to be that we can evaluate whole systems of explanation on the basis of theoretical virtue.

‘[W]hen the truth of *P* causally explains the truth of *Q*, we may still maintain that the truth of *Q* consists in something more (or other) than the truth of *P*. Or again, the fact that someone broke a promise may “normatively” account for his having done something wrong, but that is still compatible with his wrongdoings consisting in something more than his having broken the promise. There is, however, *no explanatory connection that stands to ground as grounding stands to these other forms of explanation.*’ It is the ultimate form of explanation; and it is perhaps for this reason that we are not inclined to think of the truth of a grounded proposition as a further fact over and above its grounds, even though it may be distinct from its grounds...’ (Fine 2001: 15-16, emphasis added).

So Fine takes ground to be the tightest explanatory connection between truths, but it isn’t clear what we should make of this. Central to Fine’s account seems to be that a grounding explanation is an explanation such that there is a constitutive connection between the explanandum (or the grounded fact) and the explanans (the grounding fact). One might be tempted to see this as an expression of the idea that grounded entities are an ‘ontological free lunch’, but Fine is explicit in his statement that the grounded facts might themselves be ‘real’ (2001: 16), where Fine identifies the real with the fundamental; Fine thinks that fundamentality and ungroundedness can come apart. So it seems as though we can have constitutive determination relations between fundamental facts.

Here is a worry about Fine’s attempt to clarify the notion of metaphysical explanation. Perhaps a person’s wrongdoing consists in something more than his having broken a promise (as Fine says above), and so his promise breaking does not provide the ‘ultimate explanation’ for his wrongdoing. But presumably any ultimate explanation for his wrongdoing must *involve* his promise breaking – perhaps the truth of the proposition that he did something wrong consists in the conjunction of the fact that promise breaking under any circumstances is morally wrong, and the fact that the person broke his promise. So it looks like the promise breaking is a *partial* ground for his having done something wrong, and that the promise breaking plus some other truths are full grounds for his having done something wrong. Fine must think that metaphysical explanation goes with full grounds (otherwise he would think that the

promise breaking *was* a metaphysical explanation of the wrongness). Now if metaphysical explanations are built up from a specification of other, different sorts of explanations (such as normative explanations and natural explanations), then it sounds as though (in contrast to what Fine, 2012a: 50) claims – see Chapter 1, section 2.1) full ground is to be understood in terms of partial ground. The alternative is that partial grounds are partial metaphysical explanations, but given what Fine says about metaphysical explanation (that ground is the tightest explanatory connection between truths, that we aren't inclined to think of the truth of a grounded proposition as a further fact over and above its grounds) render the notion of a partial metaphysical explanation all the more opaque.

This may seem like an *ad hominem* attack on Fine, and so to generalise worries about Fine's notion of metaphysical explanation to the notion in general may seem unjustified. The nature of my complaint, however (and the basis of the rest of this chapter) is that the notion of metaphysical explanation is opaque. Fine's account leaves me no clearer on how to understand the notion of metaphysical explanation, and raises an additional worry about how to synthesise an in-any-case-murky notion with other features of his account of grounding. Other authors offer no account of metaphysical explanation at all, or offer so little that it is very unclear what they take the notion to be.

Connecting the notion of ground to that of metaphysical explanation is supposed to shed light on the grounding relation. But as Daly (2012: 94) points out, if talk of metaphysical explanation is merely talk of grounding by another name, it cannot help to advance our understanding of either notion. In particular, the concept of grounding cannot be connected to explanatory concepts in an illuminating way simply in virtue of being labelled 'metaphysical explanation'. For the connection to be an informative one, we need an independent handle on the notion of metaphysical explanation.

In the next section, I attempt to cast light on the notion of metaphysical explanation by appeal to theories of explanation as they are discussed in the philosophy of science. There, much ink has been spilled in the quest to give an adequate account of the phenomenon of explanation. I'll argue that though it seems as though theoretical frameworks for scientific explanation could be extended to give accounts of metaphysical explanation, such accounts would then inherit the pragmatic features of explanation

which are common, to a greater or lesser degree, to all major theories of scientific explanation. In section 3 below I'll show how this, in conjunction with some other theses about grounding, leads to a problem.

2. Models of explanation

Though it is common to find philosophers helping themselves to the notion of metaphysical explanation in a discussion or defence of grounding, it is rare to find even a brief account of what metaphysical explanation amounts to (or, for that matter, of what is meant by any notion of explanation). If the explanatory connection is to shed light on grounding, an account of metaphysical explanation is owed. This must be an account of both what is explanatory about metaphysical explanation, and of what is distinctively metaphysical about it.

Since no detailed account of metaphysical explanation is forthcoming in the literature on grounding, my project in this section is to look at accounts of explanation which have been developed in the philosophy of science, and to attempt to extend and adapt these models to cover purported instances of metaphysical explanation. Discussions of explanation in the philosophy of science are usually thought to carry over to more general forms of explanation (see Woodward, 2011). For this reason, when I mention 'ordinary' explanation, I have in mind approaches based on accounts of explanation in the scientific case. Consulting extant approaches to explanation as are available in the literature should provide us with adequate frameworks for understanding explanation in the metaphysical case. In other words, my hope is that we can look to the relevant literature to give us an account of the explanatoriness of metaphysical explanation.

I (very roughly) sketch three major approaches to scientific explanation and assess how they might be adapted. The first is the covering-law model influentially advocated by Hempel; the second is the causal model advanced by Salmon (I also mention Lewis' causal account of explanation); and the third is the account of explanation associated with van Fraassen's constructive empiricism, which I call the pragmatic approach to explanation. The upshot of the discussion is that an account of metaphysical explanation that fits with our current theorising about explanation will have to incorporate some pragmatic features.

2.1. The covering-law model

The explanandum of a given explanation is the fact (or phenomenon) to be explained. The explanans is whatever explains the explanandum. The covering-law model of explanation is so-called because on that model, explanations always appeal to laws as part of the explanans. To give an explanation for some fact or event is to show that it follows from a law (or set of laws) and a specification of some initial conditions. Most often the form of this model that is discussed is the deductive-nomological (DN) account. DN explanation takes the form of a sound deductive argument. That is, ‘the explanandum must be a logical consequence of the explanans’, and ‘the sentences constituting the explanans must be true’ (Hempel, 1965: 248).

This account can fairly obviously be applied to a case of purported metaphysical explanation. Our explanandum (e.g. the existence of the set with Socrates as its sole member ($\{\text{Socrates}\}$)) must be shown to follow logically from some initial conditions (e.g. the existence of Socrates) and some laws (e.g. the set-theoretic laws).³² So far, so good. In this case it seems as though there is a sound deductive argument from some initial conditions and some laws to the relevant conclusion (the existence of $\{\text{Socrates}\}$), and therefore that the DN account of explanation can be extended to cover purported cases of metaphysical explanation. This is the first step in providing an account of metaphysical explanation (though it cannot be the whole story – we’ve not yet given a plausible account of what is distinctively *metaphysical* about metaphysical explanation, only what is explanatory about it).

However, there are some cases of purported metaphysical explanation that are not such a clear fit with the DN model. Hempel takes the aim of explanation to be understanding. Thus, the thrust of Hempel’s idea is that ‘given the particular circumstances and the laws in question, the occurrence of the phenomenon *was to be expected*, and it is in this sense that the explanation enables us to *understand why* the phenomenon occurred’ (1965: 337, italics in original). In the case of Socrates and $\{\text{Socrates}\}$, it is the case that the fact that (knowledge of) the existence of Socrates and competence with the relevant set-theoretic axioms will lead us to both expect and understand the existence of $\{\text{Socrates}\}$. But we can

³² The vast literature on distinguishing laws from accidental generalizations is of course relevant to any serious discussion of the DN account, though it is beyond the scope of this thesis.

construct examples involving purported cases of metaphysical explanation where the extension of the DN model fails to provide us with understanding of the explanandum.

Assume (as is orthodox) that grounding relations are transitive, and that when y grounds x , y metaphysically explains x . It follows that *any entity* y that is in the grounding chain for x metaphysically explains x . The extension of the DN account to metaphysical explanation dictates that (facts about) y , along with some laws, must form the explanans for x , and that the explanandum must be a logical consequence of the explanans. The risk is that when x and y are connected via a long chain of grounding, y (and some laws) is an explanation for x whether or not we, with our various cognitive limitations, can come to expect x on the basis of y and the relevant laws.

Here's an example. As described in Chapter 3 (section 3.1) Jonathan Schaffer (2010a; 2010c; 2013) argues at length for the thesis he calls *priority monism* – the view that there is only one fundamental entity, and it is the entire cosmos.³³ This single fundamental entity is 'the ground of being' (Schaffer 2010a: 35), and the reality of everything else that exists (it is the concrete things with which Schaffer is concerned) derives from the existence of the whole. Metaphysical explanation 'dangles downwards from the One' (Schaffer, 2010a: 31). Entities are related in chains of grounding that 'top out' in the single fundamental entity. Since grounding relations are transitive and asymmetric, any entity x is said to be grounded by any entity y above x in the chain of grounding that ultimately terminates in the fundamental entity.

Schaffer (at least in his 2012, see section 4 below) seems to take metaphysical explanations to obtain between facts, and to track the grounding relations that he takes to obtain between objects. For simplicity, I will assume in what follows that metaphysical explanations obtain only between facts.³⁴ Because the priority monist takes the whole cosmos to provide the ultimate metaphysical explanation for each entity within the cosmos, for any given explanandum the priority monist must be prepared to take facts about the entire cosmos (along with some laws) to be the explanans. Call the relevant fact (or set of facts) that the monist appeals to 'the cosmic fact' – 'C' for short.

³³ Henceforth, any references to 'monism' in this Chapter are about the priority view.

³⁴ I take it that not much hangs on this assumption. Anybody who thinks metaphysical explanations obtain between entities other than facts may simply substitute the facts mentioned here for the relevant entities.

Now suppose we seek a metaphysical explanation for the fact that my pet cat, Jess, exists (and has the particular nature that she does). The DN account of explanation requires that the explanandum (the existence of Jess) be a logical consequence of the explanans (C, and the laws). Presumably (assuming the truth of determinism) the explanandum *is* a logical consequence of the explanans. But for creatures with our cognitive capacities, a specification of C plus the laws of nature allows us neither to expect nor to understand Jess' existence. C is simply far too general for us to be able to infer from it and the laws of nature anything as specific as the existence and nature of Jess. Note that a commitment to monism is not needed to make this point, or any of the points below which I illustrate using priority monism (I chose it as an example only because of its simplicity). The point can be made instead with pluralist foundationalism (or any account of metaphysical structure which takes grounding to be a transitive relation, and where there is sometimes a relatively long chain or large network of grounding relations). From a specification of the relevant portion of basic atoms along with the laws of nature, we would be similarly unable to infer the existence and nature of my pet cat.

The concern is that this failure seems to go against the spirit of the DN model. Hempel takes a basic motivation for science to be a 'deep concern to *know* the world...and to *explain* and thus to *understand* the unending flow of phenomena it presents' (Hempel, 1962: 9, emphasis in original). For Hempel, understanding is the consequence or goal of explanation. If our cognitive limitations are such that we do not come to expect or to understand the explanandum on the basis of the explanans then it seems that by Hempel's lights we haven't got an adequate explanation. If metaphysical explanation is a form of DN explanation, then it seems as though C ought not to count as an explanation for the existence and nature of Jess. But if metaphysical explanations obtain whenever grounding relations do, then C *does* explain Jess' existence and nature.

A second though related problem with the DN model can be seen perhaps more effectively by shifting the discussion from grounding to truth-making. Assume for the sake of argument (and with e.g. Rodriguez-Pereyra, 2005: 20; Schaffer, 2010b) that truth-making is a relation of grounding, such that truth-makers ground the truths that they make true. Truth-maker monism is the view that 'for all worlds *w*: the one and only truth-maker at *w* is *w*' (Schaffer, 2010b: 307).

Truth-maker monism has been dismissed as ‘true but trivial’ (Molnar, 2000: 83) and ‘uninteresting’ (Armstrong, 2004: 19), on the grounds that it fails to identify which particular entity is responsible for which truth. Truth-maker monism is indiscriminate. This is closely related to the problem with citing C in order to metaphysically explain a fact such as the existence and nature of Jess; the explanation is uninformative because it is indiscriminate. It doesn’t lead us to expect or understand Jess’ existing and having the nature she does any more than it leads us to expect or understand anything else (if, indeed, it aids our understanding at all).

Schaffer (2010b: 312-3) claims that the objection (in the truth-maker monism case) that people like Molnar and Armstrong are making is in fact quite hard to pin down. He considers a number of possible ways to sharpen up the objection, and finds each of them wanting. For the purposes of applying the argument to the explanation case, the most relevant is the idea that one might object to truth-maker monism on the grounds that it cannot be that multiple truths have the same truth-maker. Schaffer argues in response that (almost) all truth-maker theories will allow that multiple truths can have the same truth-maker in certain circumstances. For example, the truth-maker for my rug being coloured and the truth-maker for its being green is presumably the rug itself in both cases. It is then unclear, according to Schaffer, what can be objectionable in the further extreme characteristic of truth-maker monism (where a truth-maker is shared not just by two or three truths, but by all truths at a world).

It seems to me that there are plausible reasons for objecting to the extreme but not the limited case, and so Schaffer’s argument is too quick. It could well be that the appropriate level of grain for truth-maker theory is such that it allows some but not all facts to have the same truth-maker. (One way to mark this kind of distinction would be to look at what the facts themselves make mention of. If the rug is the truth-maker for both the fact that the rug is green and the fact that the rug is coloured, then the truth-maker for those facts is also the subject-matter of those facts. If, on the other hand, the cosmos acts as truth-maker for each of the facts about the rug, it is at least not obvious that the truth-maker is the subject-matter of the fact, since neither fact makes obvious mention of the cosmos.)

Perhaps a way to think about this that accords with our intuitions is to say that the cosmos might be a truth-maker for the fact that the rug is green, but that it is not the *most salient* truth-maker for the relevant

fact. To cite the entire cosmos as truth-maker for the fact that the rug is green is overkill; the relevant truth-maker is the portion of the cosmos responsible for the greenness of the rug. This interpretation seems to capture the thought that truth-maker monism is not false, but it is not interesting either. It is worth making the analogy with explanation explicit here – the cosmos might provide an explanation for a given fact, but it is very often not the most salient (i.e. most relevant or informative) one. (Again, the point does not require truth-maker monism. Any account whereby an entity can be cited as a truth-maker for a fact that does not have that entity as a subject matter will prove similarly uninteresting.)

The problems discussed above for the DN model of explanation seem not to be limited to the case of metaphysical explanation, but to be a species of a general concern with the DN model. The requirement that the explanandum be a logical consequence of the explanans is importantly different from the requirement that we be capable of deriving the explanandum from the explanans, or that the explanation be informative for us. Hempel's statement of his view suggests the former requirement, but the informal motivational remarks he makes about the explanans leading us to expect the explanandum, or to understand why it occurred, fit better with the latter.³⁵

We'll need some terminology here. I'll take an explanation of a given fact or event to be anything that fulfils the requirements set out by the account of explanation we're working with. In accordance with the DN account, there is an explanation of some coin's expanding when heated that appeals to the initial conditions of the universe and the laws of nature. There is another that appeals to the fact that the coin was copper, and the law that copper expands when heated. The latter I will call a successful explanation of the expansion of the coin, because it's an explanation at the appropriate level of grain, and the one we would appeal to if asked to provide an explanation for the expansion of the coin. It is the most salient explanation of the relevant fact. Accordingly, an explanation of Jess' existence and nature that cites C and the laws of nature is indeed an explanation, but it is not a successful explanation because it does not

³⁵ The worry here is part of an even more general worry that has been raised for Hempel's account of explanation. In various circumstances, the model dictates that something is explanatory when our intuitions are to the contrary (e.g. in cases of explanatory irrelevancies). Reformulations of the view make attempts to rectify the problem, but are beyond the scope of this thesis.

lead us to expect or to understand the explanandum. With this in mind, we can proceed to discuss the causal model of explanation.

2.2. The causal model

The causal model of explanation was first properly developed in Salmon (1984). It states that explanations are explanatory only in so far as they succeed in depicting a *causal* relation between explanans and explanandum. The obvious objection to applying the causal model of explanation to the metaphysical case is that grounding is generally supposed to be a non-causal form of dependence, and so it would be *prima facie* strange to think that grounding shares a close connection with explanation, where explanation is to be understood as causal explanation. Intuitively, the existence of a singleton set depends on its member, but Socrates does not *cause* the existence of his singleton. Rather, the existence of {Socrates} is *non-causally explained* by the existence of Socrates.³⁶

Nevertheless, the heart of the causal approach to explanation can be adapted and applied to cases of metaphysical explanation. A cousin of Salmon's causal theory is the account of explanation offered by David Lewis.³⁷ Lewis (1986a: 217) takes explaining some fact to be a matter of providing some information about its causal history. An explanation thus has both a worldly component (the causal relation as it obtains in the world) and an epistemic one (the information provided about that relation). A correlate of this idea takes metaphysical explanations too to have a worldly and an epistemic component. Rather than taking the worldly component of explanation to be the causal relation (however that relation is to be understood), the proponent of metaphysical explanation can understand the worldly component of explanation to be provided by the network of grounding relations as they obtain between entities in the world (be those entities objects, facts, properties, etc.). Metaphysical explanations then provide information about this network of grounding relations.

This account of metaphysical explanation is in line with a more general trend for understanding dependence relations as the objective correlates of explanations (e.g. Kim, 1994: 67). Dependence

³⁶ Though note that assent to the non-causal nature of grounding is not universal. See A. Wilson (manuscript) for an account of grounding in terms of 'metaphysical causation'.

³⁷ The differences between the accounts are not especially important here.

relations can be causal, mereological, metaphysical, and so on, and so can the explanations that track them. This idea forms the basis of the argument that Audi (2012a: 105) advances in defence of the need for grounding relations. (See Chapter 1, section 2.2.3 for an overview and discussion of the argument).

Audi (2012a: 104) argues that there are some explanations in which the explaining fact plays no causal role with respect to the fact being explained, and so there must be some non-causal determination relation that is the objective correlate of such explanations. His examples include explaining why a sphere has the disposition to roll in terms of its being spherical. Being spherical doesn't cause the sphere to have that disposition, nor is being spherical identical with its disposition to roll. Audi claims that the fact that a given thing is spherical must therefore non-causally determine that it has the power to roll, and that this non-causal determination relation is grounding.

A model for metaphysical explanation that takes grounding relations to be the worldly component of such explanations seems promising; it gives us a model for metaphysical explanation that fits neatly into a more general story about explanation, and not only gives us an account of what is explanatory (any information about the network of grounding relations), but also offers an account of what is distinctively *metaphysical* about metaphysical explanation – it is information about the *grounding* network (rather than about any other worldly relation).

But an extension of the causal model to metaphysical explanation is not without problems. One such problem is similar in kind to that discussed above for the DN account of explanation. According to the causal model, information about *any* aspect of the causal history of an event is an explanation for that event. On our extension to metaphysical explanation, information about any aspect of the grounding chain or network that grounds the entity for which we seek an explanation is a metaphysical explanation for that entity. Built into the causal model, and correspondingly into our extension of the causal model, is the idea that explanations might be better or worse. Using the terminology introduced at the end of the previous section, we can say that any information about the grounds for the entity for which we seek an

explanation is a metaphysical explanation for that entity, but that only some of these explanations are successful.³⁸

Lewis (1986a) identifies a number of ways in which an explanation might fail to be satisfactory (or successful in my terminology). One such way is that the information provided might be ‘a true but weak proposition; one that excludes few...of the alternative ways the causal history of the explanandum might be’ (Lewis, 1986a: 226). A metaphysical explanation will be thus unsuccessful if it fails to exclude many alternative grounding stories for the entity in question. This is closely related to the constraint on the DN model that a successful explanation should lead us to expect or to understand the explanandum.

Let’s think briefly about Jess that cat and monism again. Monism tells us that the cosmos grounds Jess, and so facts about the cosmos provide a metaphysical explanation for facts about Jess. Moreover, a monist is very likely to cite facts about the cosmos when asked for a metaphysical explanation for Jess. But an explanation that appeals just to the cosmic fact seems to violate Lewis’ condition on successful explanation because it fails to exclude many alternative grounding stories for Jess, at least given our limited cognitive capacities. Again, the cosmos might provide *an* explanation for the existence and nature of Jess, but it seems not to provide a successful explanation. We cannot come to understand Jess’ existence and nature on the basis of the cosmic fact, and we cannot rule out a vast number of alternative grounding stories mediate between Jess and the cosmos.³⁹

The issue is that the metaphysical explanation offered by a monist for Jess’ existence and nature will cite the existence and nature of the cosmos, but by the lights of the causal model this explanation is unsuccessful. The problem does not depend on the truth (or even the possibility of the truth) of monism. The very same issue will arise for a pluralist account of metaphysical structure (a specification of the fundamental atoms does not allow us to exclude the possibility that this table have other entities in its grounding chain than those that that in fact ground it). If metaphysical explanations simply provide

³⁸ The terminology is not important here. I do not wish to quibble about whether or not unsatisfactory explanations deserve to be called explanations, or whether or not ‘good’ and ‘satisfactory’ and ‘successful’ are synonymous as applied to explanations. For my purposes here, I will assume that they are.

³⁹ One might object that this is to interpret the conditions on successful explanation in a way that is too epistemic, and that a more metaphysical understanding of these conditions is appropriate in the case of metaphysical explanation. My response is explained below; if metaphysical explanation is to elucidate grounding, those explanations must be recognised as explanatory by us.

information about grounding relations, then the cosmos or the atoms do provide metaphysical explanations. If these explanations are to be considered unsatisfying or even unsuccessful (as Lewis' comments imply) then we should stop appealing to them as though they were explanatory. If, on the other hand we wish to continue to appeal to such explanations, then it seems that the goal of metaphysical explanation must be something other than that of causal explanation.

In defence of the causal account of metaphysical explanation, one might wish to distinguish between explanation and *ultimate* explanation. Just as there are intermediate grounds between Jess and the cosmos, perhaps there are also intermediate metaphysical explanations. (Compare an account of what caused the window to break that cites the Big Bang with one that cites the brick that was thrown at the window.) One might suppose that this observation resolves any concerns that metaphysical explanations are sometimes apparently uninformative, and thus unsuccessful. For informativeness, we can appeal to facts closer in the explanatory chain to the fact for which we seek an explanation (without giving up on the idea that all metaphysical explanation ultimately comes from C).

To see the problem with this proposal, recall the purpose of our search for an adequate account of metaphysical explanation. What matters, if we are to be able to appeal to metaphysical explanation to get a handle on grounding, is that we are able to recognise the explanation in question. It is no use to us for our theorising about grounding if there is an explanation that corresponds to the grounding relations, but it is not the explanation we consider successful, or would offer when asked for an explanation. We require an account of metaphysical explanation that connects grounding with the explanations we recognise; those that increase our understanding of the phenomena in question.

2.3. The pragmatic approach

van Fraassen's pragmatic approach to scientific explanation takes explanations to hold relative to a theory, independently of whether that theory is itself true or even believed or accepted (van Fraassen, 1980: 101). van Fraassen is motivated by the need to draw out particular features of the causal net that science describes in order to give a satisfying explanation (a motivation with which we can sympathise, given the difficulties discussed for the causal account). The explanatory factors are to be drawn from a

range of factors that the scientific theory lists as objectively relevant, but the choice of which of these are in fact explanatory (or make for satisfying explanations) is determined by the context of the explanation request (van Fraassen, 1980: 126). The context for the explanation gets picked out by the asking of an appropriate ‘why-question’, to which the explanation provides the answer.

Here’s an example. Suppose that what requires explaining (the explanandum) is the absence of a painting from a gallery collection. The explanation for the painting’s being missing in response to the question ‘where has the painting gone?’ would be different to the explanation in response to the question ‘why would anyone steal *that* painting?’, even though responses to either question could count as explanations for the absence of the painting from the collection. An answer to the first question (‘it was stolen’) would fail to satisfy anybody who asked the second. Note that the context might differ (and so different explanations count as explanatory) even when the why-question is the same. Suppose that it was a background assumption in one context that a painting had been stolen, but not an assumption in a different context. As above, responses to a question like ‘why is that painting missing?’ will differ depending on the context of the explanation request. It is the context of the explanation request that determines which is the most satisfying explanation.

A straightforward application of the pragmatic account to metaphysical explanation takes metaphysical explanations to be answers to appropriate why-questions, which will be successful or otherwise dependent on the context of the explanation request. van Fraassen holds that an answer is explanatory for a subject if it confers greater probability on the fact to be explained (P_e), than on any other P_i in the relevant contrast class (X),⁴⁰ where the contrast class is a function of the context as determined by the question. An application of Bayes’ Rule allows one to calculate the probability of a given event relative to background conditions and some new information. van Fraassen takes the background conditions (a body of accepted background theory and factual information) to be dependent on who the questioner and the audience are (1980: 145). This will vary with context, but is held fixed in an evaluation of an

⁴⁰ Note that there are apparent examples of probability-lowering explanations in the literature (where the explanation trumps a more reliable backup), but I do not discuss them here.

answer. Therefore, what counts as a good, satisfying or successful explanation for one person might not for another.

This pragmatic account of explanation will not be subject to the problems outlined for the DN and causal models as applied to metaphysical explanation, because the pragmatic account takes explanations to be relative to a theory. What is explanatory for a priority monist might differ from what counts as a satisfying explanation for a pluralist, or for anybody with a less well-developed view of metaphysical structure.

Fine (2001: 22) takes the search for a metaphysical explanation to be an attempt to explain ‘in the most metaphysically satisfying manner’ why one proposition makes another true, and so it may seem that van Fraassen’s account is quite far removed from what we seek in an account of metaphysical explanation. However exactly we are to understand the notion of the ‘most metaphysically satisfying’ explanation, it seems at least to suggest something objective, not something determined by the context of the explanation request or by the background commitments of the audience. To reconcile the idea that metaphysical explanation is explanation in the most metaphysically satisfying manner with the idea that explanations are context sensitive, we could take the idea of a ‘metaphysically satisfying’ explanation as an attempt to fix a context – metaphysical explanation is explanation in a metaphysical context. (Compare a debate over the existence of tables, where both parties agree that there are tables, but only one accepts that ‘metaphysically speaking’ there are tables.)

For a given question, we might then find that the causal explanation and the metaphysical explanation come apart. So, for example, in answer to the question ‘why is it raining?’ a causal explanation might mention a collision between a mass of warm air and a mass of cooler air. A metaphysical explanation might focus instead on specifying the conditions under which we are entitled to apply the predicate ‘is raining’. A causal explanation of why Bill did something wrong might focus on his anger at Bob’s having slept with Bill’s wife, which led to Bill murdering Bob. A metaphysical explanation of the same fact might focus on the dependence of moral truths on facts about pleasure and pain (similar distinctions are made in Trogon, 2013b). This would give us a clear account of when a given metaphysical explanation

was metaphysical, without us having to endorse any greater degree of context sensitivity to explanation than merely determining the appropriate type of explanation as requested by the question.

Aside from the fact that limiting the pragmatic approach in this way would mean losing some of the benefits of the approach (such as the ability to cope with the sorts of difficulties raised for the DN and causal accounts), it seems implausible that we could limit the context dependent elements of the pragmatic approach to specifying a context. Specifying that the relevant context is ‘metaphysical’ will still allow for multiple seemingly successful candidate explanations of the same fact or event. A metaphysical explanation of x that explains why x obtains at all might be very different from a metaphysical explanation of x that explains why x *rather than* y obtains. For example, the metaphysical explanation of why some action is morally permissible might cite its conformity with the Kantian universalisability principle, but its being permissible *rather than required* is explained by its being a supererogatory act. Even for the very same agent, these different explanations will be the most appropriate, satisfying and successful explanation of the relevant fact in response to different why-questions, and so the context sensitive elements of van Fraassen’s account cannot be limited to specifying a metaphysical context. (For the account to work, the context sensitivity of explanation must be more fine-grained.)

2.4. Metaphysical explanation and explanation-seekers

We can learn a number of lessons from the above discussion of models of scientific explanation, and their application to the case of metaphysical explanation. First, we can note that it is fairly easy to adapt these scientific models in order to give an account of the explanatory component of metaphysical explanation (a sound deductive argument; information about the grounding network; an answer to a particular why-question). However, only one of the models discussed above offered a way to understand what is distinctively metaphysical about metaphysical explanation; the extension of the causal model has it that metaphysical explanations are information about the grounding network. The other models of scientific explanation cannot help us to identify the metaphysical aspect of metaphysical explanation (we dismissed an attempt to do this in the context of the pragmatic approach).

The second thing to note is that it is important to distinguish between explanation in an objective sense (as might be given by a sound deductive argument or a true causal statement) and explanation in an interest-relative sense. I take the above discussion to demonstrate that the notion of interest is that latter notion, because *successful* explanations (as I have understood them) belong to this latter category. A good account of metaphysical explanation must be (at least in part) an account of successful metaphysical explanation if such explanations are to be properly informative in a way that aids our understanding. In particular, if our account of metaphysical explanation is to shed light on the problematic notion of grounding, the explanations in question must be those we recognise as explanatorily successful. (If they lacked this feature, we would be no clearer on the notion of grounding because we would be trying to understand grounding via a notion of explanation that is itself murky and confused). This line of argument is explored in more detail in section 4 below.

Of the three models discussed, it is the pragmatic approach that best satisfies the role that metaphysical explanation is taken to play in relation to grounding. This model affords us an account of explanation that we can extend in such a way as to always generate a satisfactory account of metaphysical explanation that accords with our explanatory intuitions, as is required if metaphysical explanation is to shed light on grounding. The key is that the background beliefs and commitments of the explanation-seeker are built into the account of explanatory success, and so the explanations that are successful for an explanation seeker will be successful metaphysical explanations (in the given context). The approach might therefore seem to come at a high price. (Successful) explanation is a highly context-dependent phenomenon, because the success of a candidate explanation depends on features of the agent who seeks the explanation.

In fact, the pragmatic account does not differ so markedly from the other approaches in this last respect. Recall that it was part of the motivation for the DN account that explanation leads us (the explanation seekers) to *expect* the explanandum on the basis of the laws and initial conditions, and to *understand why* it occurred. Expectation and understanding are agent-relative notions. In the causal model, unsatisfying explanations are those that fail to speak to the interests of the agent receiving the explanation, or are too general to allow her to exclude alternative histories for the explanandum. To

provide a good explanation is, in part, to describe aspects of the causal history of the explanans in such a way that they increase understanding in the recipient of the explanatory information.

One might object here that there is an element of idealisation in these accounts of explanation; it would be no failing if the account was such that a two-year-old would fail to expect or to understand the explanandum on the basis of the explanans. Similarly, one might argue, our adult epistemic limitations fall short of a reason to build an element of context sensitivity into an account of explanation. So long as an ideal epistemic agent were able to expect and to understand the explanandum on the basis of the explanans, we should count the relevant explanation as successful. In response, note again that we need a notion of metaphysical explanation that is suited to elucidating grounding, and that accounts for our practice of appealing to explanatory intuitions to identify what grounds what. For these strategies to be successful, the relevant notion of metaphysical explanation must be one that produces explanations that lead not the ideal epistemic agent, but creatures like us to come to expect and understand the explanandum on the basis of the explanans.

What it is to be the most successful or best explanation of a given fact is to be the most salient explanation in a given context, and this notion is implicit in each of the accounts under discussion. The extra component to successful explanation essentially involves explanation-seekers. In what follows, I will refer to whichever features of explanation involve explanation-seekers as ‘pragmatic’ or ‘agent-relative’.

3. Articulating the problem

We are now in a position to state a problem for our theorising about grounding and metaphysical explanation, which boils down to a tension between three independently plausible theses:

- (i) There is a tight connection between grounding and metaphysical explanation
- (ii) Metaphysical explanation has pragmatic, or agent-relative features

(iii) Grounding is a feature of objective, mind-independent reality⁴¹

Both (i) and (iii) are cornerstones of the grounding literature (see section 1 of this Chapter, and section 2 of Chapter 1). The most controversial claim amongst the three theses is the second, which is defended above. The resulting problem is as follows. If grounding is a feature of objective and mind-independent reality, but explanations are relative to an agent, then it seems the connection between grounding and explanation can't be as close as the literature suggests. If the connection is to be maintained along with the objectivity of grounding, then either explanation must be similarly objective and mind-independent (contradicting (ii)) or grounding, like explanation, must have agent-relative features. Call this tension 'the grounding-explanation problem'.

Let's hold fixed that grounding is a feature of mind-independent reality. Now we have a dilemma – we must either reject the thesis that metaphysical explanation is pragmatic, or deny the connection between grounding and metaphysical explanation. If (as I will argue) neither of these options can be made palatable, we should consider rejecting the assumption with which we started. This will be the topic of the final Chapter.

4. First horn: metaphysical explanation

Let's first consider rejecting the thesis that metaphysical explanation has pragmatic features. The discussion above demonstrates attempts to extend our understanding of scientific explanation to generate theories of metaphysical explanation. This involves an implicit assumption that metaphysical explanation is in the same mould as scientific explanation, but we need not buy into this assumption (particularly if we are unhappy with where it leads us). The alternative to understanding metaphysical explanation as by extension from theories about scientific explanation is to take metaphysical explanation to be distinct from ordinary explanation.

⁴¹ In Chapter 5 I discuss problems with articulating a conception of mind-independent reality. For now I rely on an intuitive distinction between the mind-dependent and the mind-independent, such that the latter but not the former exists independently of linguistic and conceptual schemes.

It would be open to anybody who adopts this view of metaphysical explanation as distinct to deny that metaphysical explanation inherits the pragmatic features of ordinary or scientific explanation. If metaphysical explanation does not have pragmatic features, the grounding-explanation problem dissolves. Below I explain why I think such an attempt to dissolve the problem will be unsuccessful. I then discuss the possibility of restricting the agent-relative features of explanation to *successful* explanation, and argue that the relevant connection between grounding and metaphysical explanation is between grounding and successful metaphysical explanation. Finally, I consider appraising systems of metaphysical explanation by appeal to theoretical virtue rather than to explanatoriness for agents, and argue that there is no objective way to weigh up theoretical virtues.

First, we should note that in order to maintain that ordinary explanation and metaphysical explanation are distinct to such an extent that ordinary explanation is sensitive to pragmatic features where metaphysical explanation is not, the defender of such a view ought to be able to distinguish between the two types of explanation. But it is not obvious that this is always possible. deRossett (2013b: 243) claims that ‘grounding explanations are exemplified by certain familiar scientific explanations’, and cites a number of examples. An isotope of gold has its atomic mass in virtue of containing a certain number of protons and neutrons. Diamond is hard because each carbon atom in the structure is bonded to each of its neighbours. We might add that ice floats on water in virtue of the hydrogen bonds in ice crystals that space the molecules further apart than when they are in a liquid, making ice less dense than water. Or take a non-scientific example, such as an explanation of the success of a birthday party that cites the number of people present. In each of these cases we have an explanation that is both metaphysical *and* either scientific or ‘ordinary’. It is implausible then that scientific explanation is sensitive to pragmatic factors, but metaphysical explanation is not.

Let’s assume for the sake of argument that anybody wishing to maintain a clear distinction between ordinary and metaphysical explanation has an answer to this sort of concern, and grant that metaphysical explanation is entirely independent of any context or agent. Whoever adopts this dualistic position with regards to explanation faces unpalatable consequences. It is the aspect of metaphysical explanation which is akin to ordinary explanation (e.g. intuitive understanding of terms such as ‘because’ and ‘in virtue of’;

formal features such as asymmetry, transitivity, non-monotonicity and hyperintensionality; and a sense of what it is to convey answers to why questions) that allows us our intuitive grasp on metaphysical explanation in the first place.

Success in arguing that metaphysical explanation is distinct from ordinary forms of explanation renders it difficult to give an account of what metaphysical explanation actually *is*, and thus very hard to see how a connection with metaphysical explanation could shed light on grounding. The connection between grounding and metaphysical explanation is only illuminating in so far as the relevant form of explanation is similar enough to our ordinary conception of explanation for us to understand it. To deny that metaphysical explanation is sensitive to context (as is ordinary or scientific explanation) is to cast that notion into obscurity. In sum, arguing that metaphysical explanation is unlike ordinary explanation in such a way that the former lacks the pragmatic features of the latter is difficult because in some cases it will be hard (or even impossible) to tease the two apart, and undesirable because it undermines the utility of connecting the notion of grounding to that of (metaphysical) explanation in the first place.

We are searching for ways to defuse the grounding-explanation problem that involve rejecting the thesis that metaphysical explanation has pragmatic features. Forcing a robust difference between ordinary (scientific) and metaphysical explanation with respect to pragmatics does not exhaust our options. Instead, we might argue that these pragmatic features are not really affecting what it is for something to be an explanation, but only what it is for something to be a *successful* explanation.

At least according to the covering-law model and the causal theory of explanation, the supposed pragmatic elements of successful explanation do not really affect what makes for an explanation. Instead, they merely affect what is the most salient or successful explanation of a given fact or event, and this is what differs with context. We can explain an event by citing any or all of its causal history, or a fact by citing any number of laws and initial conditions from which it can be derived. A specification of these features might not count as a successful or desirable explanation for a particular agent in a given context, but they are nevertheless *explanations*. Perhaps we can get by with connecting grounding to a bare notion of metaphysical explanation without getting into the issue of what makes for successful explanation, and thereby avoid the problematic pragmatics.

Here is one seemingly attractive way such a proposal might go: Just as causal explanations cite portions of the causal history of an event, metaphysical explanations cite portions of the grounding story for a given fact (or entity). This proposal only becomes problematic when we need to appeal to metaphysical explanation in order to elucidate grounding, but that is precisely what our project is here. Recall the ways in which the connection between grounding and explanation has been defended: explanatory dependence is more fine-grained than modal dependence, explanatory locutions in ordinary language are grounding locutions, grounding and explanation share formal features, and appeal to intuitions about what explains what can tell us what grounds what.

Assume that the relevant notion of metaphysical explanation does not have pragmatic features. This does not affect the claim that explanatory dependence is more fine-grained than modal dependence, or that grounding and explanation share formal features. However, it does reduce the plausibility of the claim that explanatory locutions are grounding locutions, and causes serious difficulties for the idea that we can appeal to intuitions about what (metaphysically) explains what to tell us what grounds what.

When people make claims using locutions that are purportedly grounding locutions, they do so by way of explaining something to somebody (e.g. why a certain act is wrong). If all our evidence of folk uses of grounding locutions involve claims about what we have been describing as successful metaphysical explanation, rather than claims about ‘mere’ metaphysical explanation, we have reason to believe that the notion of successful metaphysical explanation is the important one. An elucidation of grounding that connects grounding locutions with explanatory aspects of folk discourse connects grounding with the pragmatically laden notion of *successful* metaphysical explanation.

Of course, one might respond that the folk have no need for a context-independent notion of metaphysical explanation, just as they have no need for a notion of causal explanation that doesn’t build in relevance. Just as they report only the salient features of the causal network, they report only the salient features of the grounding network. We don’t doubt that there can be a notion of causal explanation that doesn’t build in agent-relativity; it is the model that says that information about any aspect of the causal history of an event is a causal explanation. We should be correspondingly happy to

accept that the important notion here is metaphysical explanation and not successful metaphysical explanation, despite what folk use of grounding locutions seems to suggest.

In response, note that we are not using a folk concept of explanation-talk in order to get a handle on causation, and there is no suggestion that causation itself is an explanatory notion. Instead, we have a number of candidates for understanding the causal relation (via counterfactual analyses, statistical theories, causal realism, and so on) none of which appeal to the notion of causal explanation. The close connection between grounding and metaphysical explanation is what makes a connection between grounding and a form of explanation not found in folk discourse hard to swallow, especially if we are to appeal to folk discourse to argue that grounding claims are intelligible.

Where a lack of pragmatic elements of grounding claims would do the most damage to the grounding theorist is in the way in which intuitions about metaphysical explanation are thought to provide evidence for grounding claims. Intuitions about what explains what are intuitions about *successful* explanation, and so we are only able to appeal to intuitions about successful metaphysical explanation to tell us what grounds what. It is thus far more plausible that the notion of ground is to be connected with a notion of successful (and thus pragmatic) metaphysical explanation, than with a notion of metaphysical explanation free from pragmatic features.

If grounding shares a close connection with a non-pragmatic notion of metaphysical explanation, then it is a mistake to appeal to intuitions about explanation to evaluate grounding claims. As noted, explanatory intuitions concern successful or satisfying explanations. If grounding relations obtain when the corresponding metaphysical explanation is unsatisfying or unsuccessful, then appeal to explanatory intuitions offers very little evidence about what does and doesn't ground what.

There is one further option for somebody who wishes to deny that metaphysical explanation has pragmatic features. Fine (2001: 22) claims that systems of metaphysical explanation, like other explanatory schemes, can be appraised on the basis of considerations of theoretical virtues such as simplicity, breadth, coherence and explanatory strength. Evaluating these systems of metaphysical explanation is a process of uncovering the best explanation of some explanandum, and if that process

could be seen to be pragmatics-free, then one might argue that metaphysical explanation is not context-dependent or agent-relative in any interesting sense.

According to Fine, the most important of the virtues is explanatory strength; ‘the capacity to explain that which stands in need of explanation and would otherwise be left unexplained’ (Fine, 2001: 22). But in order to judge the strength of a theory, we will then be required to ascertain which things stand in need of explanation. Doing so must involve a consideration of our own requirements on explanation, but that means that once again, explanation is shown to be relative to a context – successful explanations are answers to why-questions posited by an agent with a particular set of interests, background beliefs and theoretical commitments.

Moreover, it is far from clear that judgements of theoretical virtue are mind-independent. Agents with different conceptual schemes will judge different theories as virtuous, and so appeal to virtues other than strength cannot rescue metaphysical explanation from pragmatic considerations. (See Chapter 5, section 2.2 for a more detailed discussion and defence of this line of argument.) What we take to stand in need of explanation, as well as the degree to which we judge particular explanations to be successful, will depend on pragmatic factors. Denying that successful explanation has pragmatic features is another dead end.

To recap: it is hard to preserve a distinction between metaphysical explanation and other forms of explanation, and if such a distinction were to be maintained, we would find ourselves once again without any real account of metaphysical explanation. Consequently, we would be without the independent handle on grounding that the connection between grounding and metaphysical explanation has been thought to provide. Neither restricting the pragmatic features of metaphysical explanation to a notion of successful explanation not to be connected to grounding, nor denying that a notion of successful metaphysical explanation has pragmatic features offers a fertile alternative. To deny that metaphysical explanation has pragmatic features does not make for an attractive response to the grounding-explanation problem.

5. Second horn: the grounding-explanation connection

It is unappealing to deny that metaphysical explanation has some pragmatic features, and so to resolve the grounding-explanation problem we must now consider re-evaluating the purported link between grounding and (successful) metaphysical explanation. I take it that there are broadly three ways to understand the connection. First, we can think of the explanatory role of a particular instance of grounding as an entirely *extrinsic* feature of that relation, so that grounding relations have explanatory value relative to particular explanatory projects, but are not explanatory in themselves. Second, we might think that grounding relations have a closer connection to explanation such that the grounding relation *just is* a relation of explanation. Both accounts of the grounding-explanation connection can find support in the literature. Alternatively, we can deny that grounding and explanation are closely connected at all. I'll return to this third option at the end of the section.

Take the first view. Audi (2012a) insists that 'grounding is not a form of explanation, even though it is intimately connected with explanation (119)... an explanation...is something you can literally *know*; a grounding relation is something you can merely know *about* (120)', and thus seems to take metaphysical explanation to be something outside of the grounding relations. A metaphysical explanation might be a proposition expressing a grounding relation, but is not the relation itself. Schaffer (2012: 124) advises 'one should distinguish the worldly relation of grounding from the metaphysical explanations between facts that it backs, just as one should distinguish the worldly relation of causation from the causal explanations between facts that it backs'. Schaffer seems to assume here that explanations are facts involving the relata of grounding relations, but that the grounding relations themselves are not explanations.

It is hard to reconcile this way of understanding the link between grounding and explanation with the picture of grounding painted at the beginning of the paper. It is not obvious why modal notions should be too coarse-grained to capture the relevant sense of dependence if it were not explanatory, or why grounding claims are expressed with explanatory locutions in folk discourse. Nor is it obvious that we should look to the formal features of explanation to give us an account of those of grounding – after all,

we don't usually assume we can appeal to the formal features of explanation to give us an account of those of causation.

We rely on intuitions about what metaphysically explains what in order to tell us what grounds what. Furthermore, we rely on what we know about explanation (formal features and so on) to give us a characterisation of the grounding relation. If metaphysical explanation merely tracks grounding relations, then our reliance on features of explanation for our understanding of grounding is illegitimate. There is far more scope for error in our judgements about grounding than we ought to consider acceptable. There is an epistemic gap between judgements of ground and judgements of explanation that would not be present if grounding just were a form of explanation, and this gap is especially troubling given the severely limited supply of alternative methods for discovering facts about grounding. This conception of metaphysical explanation robs us of the special insight into grounding that metaphysical explanations have been supposed to provide.

In response to such concerns, a defender of the extrinsic conception of the grounding-explanation link might stress that the connection, though extrinsic, is very tight. Perhaps there is something like a synthetic necessary connection between grounding and metaphysical explanation, so that whenever there is a grounding relation there will be a metaphysical explanation that corresponds to it. The problem (besides any concerns about necessary connections) is that in tightening the link, a defender of the tight extrinsic conception opens herself up to other problems that she faces in common with a defender of the view that the connection between grounding and explanation is *intrinsic*.

Support for the intrinsic conception of the link can be found in a number of Fine's works. Fine claims that 'the relationship of ground is a form of explanation' (2001: 22) and that

'[it] is properly implied by the statement of (metaphysical) ground...that there is no stricter or fuller account of that in virtue of which the explanandum holds. If there is a gap between the grounds and what is grounded, then it is not an explanatory gap.' (Fine, 2012a: 39.)

In a similar vein, Rosen says that ‘the grounding relation is an explanatory relation’ (2010: 117), and refers to grounding claims as explanatory claims. In much of the literature, the way in the connection between grounding and metaphysical explanation is presented is brief enough as to be indeterminate between the intrinsic and a tight extrinsic conception of the relation between grounding and explanation.

If the connection between grounding and explanation is such that grounding just is a relation of metaphysical explanation, then grounding, like metaphysical explanation, is sensitive to context. Just as an explanation of some fact is sensitive to the interests of the agent seeking the explanation, the grounding relations that obtain between those facts (or the relevant entities) are also context sensitive. This is in tension with the claim that grounding relations are a feature of mind-independent reality.

Even if grounding is not in fact a relation of metaphysical explanation but the two share a tight connection, an epistemic problem remains. When we make grounding judgements, or make claims about the grounding relation, we have no way of being sure which of our judgements are affected by the pragmatic features of metaphysical explanation, and which correspond to the mind-independent notion of grounding. Our resources for making judgements of ground are so limited that we do not have the capacity to distinguish the two notions.

Here’s an example. Take the term ‘near-sphere’ to refer to a range of objects with a more-or-less spherical shape.⁴² Amy, who is an engineer, thinks that a particular object *O* is a near-sphere because it is a dented sphere. Consequently, she takes the fact that *O* has a dent to ground the fact that *O* has the (maximally determinate) shape *S*. Barry, who is an artist specialising in drawing abstract shapes, thinks that the fact that *O* has shape *S* grounds the fact that it is more-or-less spherical. It is implausible that both of these grounding claims are true, as then we would have to conclude (assuming grounding is transitive) that the fact that *O* has a dent grounds the fact that it is more-or-less spherical. As the presence of the dent seems irrelevant to the sphericity or otherwise of *O* (it could have been a dented triangle), we ought not to accept the claim.

⁴² This example is adapted from Schaffer (2012: 126).

Each person's judgement about grounding is influenced by their judgement about which fact is explanatory, and which fact they take to be explanatory differs in accordance with their differing background beliefs and commitments. Amy's and Barry's judgements of ground are infected, without their knowledge, by differences in context (where the relevant context here is the contrast class of alternatives).⁴³ Without the benefit of assessing judgements of ground from the perspective of one who knows the relevant features of the context (as is our position with respect to these toy cases), it is impossible to separate out context-sensitive elements of grounding judgements. And even in cases where there is near-universal agreement about what grounds what (e.g. that Socrates' singleton is grounded in Socrates) there is no guarantee that our judgements of ground are not relative to a context shared by all those involved in discussions of grounding (but not shared, for example, by creatures with different experiences or interests, or with different conceptual schemes).

A related problem applies when we attempt to characterise the grounding relation itself. We might assume that there are some features of explanation which are also features of grounding, when in fact those are not features of the objective, mind-independent grounding relation. By way of an example, consider the contention that grounding is asymmetric. The asymmetry of grounding is thought to be a species of the general asymmetry of explanation (see Lowe, 2010). The asymmetry of explanation is a consequence of the thought that symmetric explanations are bad explanations because they fail to provide us with any new information. As I argued in Chapter 3, some cases of symmetrical grounding ought to be considered quite plausible.⁴⁴ If it were true that symmetrical explanations are ruled out on pragmatic grounds, but there *are* cases of symmetrical grounding, then the epistemic features of explanation would be influencing our theorising about grounding in a way that led us to believe false things about the grounding relation (i.e. that it is symmetric rather than nonsymmetric).

⁴³ This particular problem can be solved by moving to a contrastive account of grounding (as developed in Schaffer, 2012) whereby it is the fact that x rather than y that grounds the fact that x' rather than y' . More complicated differences in background commitments of agents making judgements of ground are not obviously resolved by moving to a contrastive account.

⁴⁴ As it happens, my view here is more nuanced. I think that *both* grounding relations and explanations can be symmetric in the right contexts, but arguments for symmetric grounding do not have to depend on the idea that explanation is sometimes symmetric.

All of this suggests the possibility of severing the link between grounding and explanation, and maintaining that the two concepts are independent and distinct (though they may of course be similar in various respects). The upshot of this would be that it would not be safe to infer anything about the character of grounding relations from that of explanation, and that any appeal to our intuitive sense of what is explanatory in a given context ought to lend at best limited credence to our judgements about grounding.

The major problem with this strategy is that it robs us of one of the most prominent approaches to making sense of the grounding relation. As we are now well aware, a crucial way to understand grounding is via its connection to explanation. If we can no longer appeal to the notion of metaphysical explanation to elucidate that of grounding, the latter is once more rendered obscure. Without the grounding-explanation link, our understanding of grounding must come from other sources that do not make appeal to the explanatory character of grounding.

This is work that has in fact been done, notably by Rosen (2010) who develops a number of general principles that involve grounding and assesses how well those principles interact with better-established metaphysical principles (see Chapter 1, section 2.2). Whilst this is a promising project, it is open to the grounding sceptic to retort that Rosen has merely defined a notion into existence that can unify some existing philosophical ideas and allow us to make some more general claims than we would otherwise be able to. This is all very well if our task is to systematise the world in a way that makes it easier for us to interpret, but not if we aim to represent it veridically. There is no reason to suppose that a notion that connects with more familiar notions in a way we find useful is therefore a part of the metaphysical manifold.

Recall that a different worry about appealing to connections between grounding and other notions in an attempt to make grounding intelligible is discussed in Daly (2012: 91-2). Daly argues that the other terms appealed to elucidate grounding are either so close to it as to be obscure themselves, or so far away that their connection to grounding is questionable (though they might be clear enough themselves). Examples of the former kind include connections to notions like fundamentality. Examples of the latter include connections to reduction, and to explanation. In the case of reduction, Daly argues that it is

standard in the literature on non-eliminative reduction to take a reduction of the *x*s to the *y*s to be the discovery that the *x*s are identical to the *y*s. Since identity is reflexive and symmetric and grounding is thought to be irreflexive and asymmetric, our prior understanding of reduction can't be used to help us understand grounding.

The method of appealing to intuitive examples of grounding to ward off the sceptic is also problematic. I explained previously (Chapter 1, section 2.2) that J. Wilson (manuscript) argues that there is no distinctive relation of (what she calls) big-G 'Grounding'. Her view is that there is a plethora of small-g 'grounding' relations – relations of type and token identity, part-whole relations, functional realisation, etc., but that these are not determinates of a common determinable. Wilson argues that philosophers almost never make general grounding claims without a more specific relation in mind (manuscript: 9). When people say that the dispositions of a thing are grounded in its categorical features, they have in mind either a token-identity theory, or a functionalist theory, and so on. Because people always have finer grained notions than Grounding in mind, there is no need to talk in terms of a more general notion of Grounding.

Finally, recall the argument against taking intuitions as evidence for grounding sketched above. When a specific example of grounding is given, we are expected to consult our intuitions in order to verify that there is a grounding relation at play here, but if intuitions about grounding are not automatically to be trusted then doubt is cast on the method of appeal to examples. Moreover, as Daly (2012: 95) argues, it is open to anybody who claims not to understand the grounding idiom to maintain that they do understand the purported examples either. If the notion of grounding is incomprehensible, so are the examples that appeal to it.

Denying the link between grounding and explanation is possible, but it is undesirable. Not only does it make it harder for us to make judgements about what grounds what, it also opens the notion of grounding up to renewed sceptical attacks regarding its intelligibility. Holding fixed the thesis that grounding relations are objective and mind-independent, a dilemma presents itself. To resolve the grounding-explanation problem, we must either deny that metaphysical explanation has pragmatic features, or find an unproblematic way to interpret the link between grounding and metaphysical

explanation. Since neither of these attempted resolutions has proven successful, we should reject the assumption with which we began – that grounding is an objective, mind-independent relation.

6. Conclusion

Not much attention has been paid to the nature of metaphysical explanation in the literature. Appeals to metaphysical explanation are made in order to shed light on the notion of grounding, but investigation into the nature of metaphysical explanation uncovers the grounding-explanation problem – a tension between our attitude to explanation, realism about grounding, and the purported link between grounding and explanation.⁴⁵ In an attempt to resolve the problem, we held fixed the thesis that grounding is mind independent and grappled with the resultant dilemma. Since we found no way to blunt either horn of that dilemma, the only path left open is to reject the assumption of the mind-independence of grounding with which we began. Antirealism about grounding offers a promising resolution to the grounding-explanation problem, and it is developed in the next chapter.

⁴⁵ Note an interesting parallel between the problems that arise for the different candidate accounts of metaphysical explanation, and the ways we might defuse the main tension. The covering-law account seems to force a disconnection between explanation and metaphysical explanation, and the causal account between grounding and explanation. The pragmatic approach allows us to retain both connections for the price of giving up on any objective stance towards explanation.

Chapter 5: Antirealism about Grounding

The main focus of this chapter is the development of antirealism about grounding. After making some comments about predominant attitudes to grounding claims in the first section, in section 2 I consider motivations for taking an antirealist perspective on grounding. In section 3 I describe some of the most influential ways in which people have defended antirealist positions as I understand them, and briefly outline how one might apply each position to the case of grounding. The most promising versions of antirealism as applied to grounding are expressivism and fictionalism, and so in sections 4 and 5 I develop models of expressivism and fictionalism about grounding in more detail, describing and outlining response to some possible objections to those accounts of grounding. I make some comparative remarks about the different accounts in section 6, and in section 7 I give a more detailed account of one motivation for appealing to grounding; to distinguish between realists and sophisticated antirealists in a given domain of discourse. I explain how antirealism about grounding is still able to retain what is attractive about this appeal to grounding.

1. Attitudes to grounding

Grounding claims are overwhelmingly made in a realist spirit. When we say, for example, that the singleton set {Socrates} is grounded in its only member – Socrates – we make a claim about what reality is like. More specifically, we make a claim about reality's structure. Reality is such that there is a connection of metaphysical dependence between Socrates and {Socrates}, such that the latter depends on, and/or is metaphysically explained by, the former.

For grounding claims to be a part of reality it need not be that the grounding entity (in this case Socrates) is metaphysically fundamental. A pluralist foundationalist might hold, for example, that the only fundamentalia are mereological atoms, and so maintain that since Socrates is a composite object he

cannot be fundamental. Since, however, Socrates is grounded in the mereological atoms that compose him (which are part of fundamental reality) and grounding is transitive, {Socrates} is firmly grounded in fundamental reality.

Recall that the recent surge of interest in grounding has been met with some resistance (see Chapter 1, section 2.2). Some are sceptical about grounding and the role that it has come to play, and this scepticism takes three main forms (see Clark and Liggins, 2012: 817-8). First, there is concern that there is no well-defined content to the notion of grounding. This arises because grounding is not given an explicit definition, but instead is introduced by example, or by citing the principles that govern it, or by the role it plays. Grounding claims are thus said to be unintelligible (see Daly, 2012). Second, the sceptic might find the notion of grounding intelligible, but be reluctant to endorse the realist interpretation of grounding claims. This could be for epistemic reasons – perhaps we could not be in a position to know about actual instances of grounding (this is a line of scepticism less often encountered in the literature, but towards which I am particularly attracted). Finally, a grounding sceptic might take grounding to be intelligible and think that there are actual instances of grounding, but think the current level of philosophical attention awarded to the notion is unwarranted. This could be because purported instances of grounding are better understood as instances of more familiar relations such as supervenience or conceptual priority. (This sort of strategy can be found in Hofweber, 2009, and J. Wilson, forthcoming).

In the last two chapters, I presented arguments that challenge the orthodox conception of grounding. Here I further challenge that view in a way that accommodates my arguments of the last two chapters. I motivate and outline *antirealist* approaches to grounding. The antirealism about grounding I envisage carves out something of a middle way between scepticism and realism, and has quite a lot in common with the second sceptical strategy mentioned above. I argue (in conflict with the first and third sceptical strategies) that grounding claims are intelligible, useful, and informative. They are not, however, claims about (or not best interpreted as claims about) the objective, mind-independent reality of the realist. Just as antirealists about mathematics or morality think that mathematical or moral claims are not really features of the mind-independent world, the antirealist about grounding denies that grounding claims

truly represent mind-independent reality. There are multiple ways in which one might be an antirealist about grounding, with positions mapping roughly onto more familiar antirealisms in other areas.

I take antirealism about grounding to be consistent with taking ground to be a nonsymmetric relation (I explain this in some more detail when I discuss specific versions of antirealism about grounding), and I think antirealism about grounding provides the most promising solution to the grounding-explanation problem discussed in the last chapter. That problem was a tension between three independently plausible theses (that grounding and metaphysical explanation are closely connected; that metaphysical explanation has some pragmatic features; and that grounding is a feature of mind-independent reality). Antirealism about grounding involves a rejection of the last of the three theses. There is no tension between the claims that grounding and metaphysical explanation share a close connection and that metaphysical explanation has pragmatic features, so long as grounding is not entirely mind-independent.

My use of the term ‘antirealism’ in this chapter differs from some ways in which the term has been understood. For example, it differs from the Dummettian proposal that antirealism is the view that the meanings of the statements in question cannot be determined in terms of (evidentially unconstrained) truth conditions (cf. Wright, 2003: 15). I use the term ‘antirealism’ here to refer to a class of views that Wright has called ‘irrealist’. In Wright’s words, ‘what opposes irrealism with respect to a particular class of statements is the view that the world is furnished to play the part in the determination of their truth values which the platitude calls for, that there really are states of affairs of the appropriate species’ (2003: 19). For my purposes then, any view whereby statements in the relevant discourse are not made true or false by a feature of the world that corresponds to the kinds of things that the statements purport to make reference to is an antirealist view. The antirealist views about grounding I am primarily interested in are those that subscribe to an additional thesis that there are some standards that determine when a grounding claim is appropriate, and that those claims *are* sometimes appropriate.

2. Motivations for antirealism

Before getting into the details of motivations for antirealism about grounding, it is worth highlighting what such a position might hope to achieve. Aside from the vocal minority of grounding sceptics, there

is a widespread sense of unease surrounding grounding and related notions. One often encounters (more often in person than in print) a feeling amongst philosophers that there is *something* to talk of fundamentality, grounding, ontological dependence, and related notions, but this feeling comes alongside a reluctance to fully endorse anything like a primitive grounding relation. This sociological observation might perhaps be explained by a concern that primitives are theoretically costly – we ought not to admit them into our theory unless we are sure we cannot do without them. Perhaps there is a conservative tendency amongst philosophers to resist relatively new bits of ideology as far as possible, or perhaps people aren't yet convinced that they understand what people are talking about when they talk about grounding. Or maybe people are (perhaps tacitly) influenced by the sorts of arguments I'll make below about grounding. In any case, I hope that antirealism about grounding might prove relatively attractive to those who hold this sort of view.

Antirealism about grounding reaps the benefits of grounding talk and suffers few of the costs. Various worries about grounding might be dampened if what grounding talk committed us to was not some primitive piece of ideology, but instead something that could be understood in other terms (as an expression of a certain kind of attitude, or a representation of another kind of content). The hope for an antirealist theory of grounding is that it might tread a delicate line between uneasiness about the notion of ground, and its usefulness in philosophical theorising. With that in mind, I'll discuss two major (and related) motivations for antirealism about grounding, both of which focus on how judgements about ground (both about local instances of grounding, and about grounding itself) are to be made. The first is that judgements of ground are made by appeal to intuitions, which seem better suited to reveal aspects of our conceptual scheme than the structure of reality. The second is that there is no objective, mind-independent method for evaluating systems of ground. Further motivations for antirealism will become apparent when specific accounts of antirealism about grounding have been introduced in sections 4 and 5.

2.1. Intuition

In this section I'll argue that the faculty of intuition falls short of providing a mechanism by which we can track grounding relations in the world. Instead, we ought to take intuitions about grounding to

provide evidence about (the structure of) our conceptual or epistemic schemes. I do not have space here to consider replies on behalf of those who take intuitions to provide good evidence for metaphysical claims, but I hope what I do say is enough to cast some doubt on the idea that intuitions about grounding reveal features of mind-independent reality.

We cannot appeal to empirical evidence to establish the truth or falsity of particular grounding claims – no amount of empirical work will tell us the answer to the Euthyphro dilemma, or whether Socrates' singleton set is grounded in Socrates or not. Moreover, it is hard to imagine that the absence of any grounding relations at all would make the world different in any way we would be aware of. But this observation is in many ways a poor motivation for antirealism about grounding, because it threatens to prove far too much. Almost all debates in metaphysics are similarly insensitive to empirical investigation (though we should not assume that all are: whilst we can make little headway on the debate between a Humean and a non-Humean about causation by conducting scientific experiments, we are at least able to make some progress in discovering what in fact causes what). There is no place for science in establishing anything either about local grounding relations, or about the nature of grounding. Consequently, we ought to pay close attention to the methods by which we do seek to establish both local grounding claims, and facts about grounding itself (e.g. its logical and structural properties).

Judgements about ground are typically made by appeal to intuition and explanation (Fine, 2001: 22). The conventional view is that metaphysicians are in the business of making claims that draw support from intuition. There are, however, reasons to be sceptical in general about the extent to which intuitions about a given metaphysical concept can tell us about mind-independent reality. The debate about the kind of justification that can be afforded by intuitions rages on, but it is clear that one way in which intuitions can provide good evidence is when they are appealed to in order to elucidate a concept (as in conceptual analysis). But conceptual analysis clearly concerns concepts, which are presumably not wholly what constitutes the realists' objective reality.

It is worth noting that intuitions about grounding are not entirely consistent. Whilst there are some cases on which there is general agreement (e.g. that sets are grounded in their members) there are other cases (e.g. grounding between parts and wholes) where answers to questions about what grounds what

are disputed. Whilst it is common to accept that wholes are grounded in their parts, as we have seen (in Chapter 3) Schaffer has argued that in the case of integrated wholes (such as circles), the parts depend on the whole (so that semi-circles are grounded in the circle). In the case of the Euthyphro dilemma too, intuitions diverge. I have appealed (in Chapter 3) to such cases in order to support my claim that grounding is in fact nonsymmetric, but anybody who denies this must then accept that it cannot be the case that everybody's intuitions are accurately representing reality. This should give us pause to wonder whether any intuitions about grounding are doing so.⁴⁶

It would, of course, be possible that though intuitions about grounding are generally reliable, some people are mistaken. This possibility would not provide a serious challenge to realism about grounding. I wish to suggest, however, that even consistent shared intuitions are unable to justify such realism. Strawson (1959) distinguishes between descriptive metaphysics – that concerned with what reality would be like if it accurately mirrored our conceptual scheme, and revisionary metaphysics – the metaphysics we would have if we had the right conceptual scheme, the one that ‘carves nature at the joints’. Intuitions about specific cases in metaphysics look like they are well placed to elucidate features of our conceptual scheme – consulting our intuitions is a good way to do descriptive metaphysics. To appeal to intuitions about cases in an attempt to find out about the real structure of reality – to help us do revisionary metaphysics, is to place our metaphysics on shaky epistemic ground.

What we can reliably expect to learn from reflecting on our intuitions about purported examples of local grounding relations (such as the relation between Socrates and his singleton) is how the entities concerned are related within our conceptual scheme, and we do not have good reasons to think that our conceptual scheme (which is affected by our theoretical commitments) provides a perfect reflection of reality. As David Wallace (2010: 69) quips, ‘our intuitions...were designed to aid our ancestors on the savannahs of Africa, and the universe is not obliged to conform to them’. It is certainly conceivable that the structure of the world could have been the same, and our beliefs about it have been very different.

⁴⁶ Even those who accept that grounding is nonsymmetric might worry that the intuitions supporting their claim are unreliable. This is why I placed little emphasis on the argument from intuition in Chapter 3. (My own antirealist view is able to draw support from divergent intuitions about the direction of particular grounding relations both in order to argue for nonsymmetry and to argue for antirealism. This will become clearer in chapters 4 and 5, where the specifics of antirealist approaches to grounding are discussed.)

This would suggest that our intuitions about the world are (at least to some extent) contingent, though its metaphysical nature is presumably not.

Owing to the richness and complexity of the debate over the nature and status of intuitions, these arguments shouldn't be understood as attempts to undermine the role of intuitions in uncovering facts about the world. My more modest intention here is to call into question the idea that intuitions about grounding provide evidence for facts about mind-independent grounding relations rather than revealing something about the way we understand the world. Whether or not our claims about grounding accurately reflect the way that the world is, the nature of the evidence for those claims (intuition) is such that we are on safer epistemic footing if we do not assume that our claims about grounding are justified by facts about mind-independent reality.⁴⁷ In section 4 and 5 I will suggest a number of alternative ways we might account for our intuitions about grounding.

2.2. Explanation

Recall that intuitions are not supposed to provide the only justification for judgements of ground. The other major source of justification for such judgements is explanatory character. Fine (2001: 22) argues that systems of ground can be appraised in much the same way as other systems of explanation, by appeal to theoretical virtue. Whatever provides the simplest, strongest, broadest, coherent, non-circular system of metaphysical explanation is the best system of ground. I'll argue that appeal to the relative virtuousness of different systems of explanation is not a reliable guide to objective, mind-independent systems of ground.

Fine asserts that the most important virtue is explanatory strength – 'the capacity to explain that which stands in need of explanation and would otherwise be left unexplained' (Fine, 2001: 22). The problem with this is that it is very hard to identify what needs explaining in any kind of metaphysical sense of the term 'explain'. That something needs explaining leaves open an argument place for context; something needs explaining to somebody, and what it is that needs explaining depends on the explanations that are

⁴⁷ I cannot engage here with works such as Williamson (2007) which hold that the evidence we get from thought experiments and the like is firmly metaphysical in character, though I acknowledge that anybody with such views will not find my arguments here persuasive. In lieu of defending my claims more thoroughly, I can point out that the argument from intuition is not my sole motivation for antirealism about grounding.

already available to the explanation-seeker (see Chapter 4, especially section 2). So perhaps Fine's use of the term 'needs' is more metaphorical than I have supposed. We might think that theories 'need' to explain the relevant data, in the sense that being a successful theory involves offering a complete (or as close as possible to complete) explanation, and that this is a sense of 'need' not obviously relative to any agent or agents. I'll work on this assumption from now on.

For Fine, the relevant data that need explaining are certain truths, and necessary connections between propositions (Fine, 2001: 22). But *which* are the truths that need explaining? This question will receive different answers from different agents, determined once again by which truths they take to be already explained, or not to require explaining for some other reason. (Recall the general worry: there is no way to evaluate judgements of ground from a standpoint that does not build in the interests of any particular agent).

Fine hopes to evaluate systems of ground by appealing to theoretical virtue, but it is open to anybody sceptical of his approach to claim that such virtues are merely pragmatic (they recommend the theory as a useful one), and that they do not add support in favour of its truth (where truth is taken to be a matter of correspondence with objective reality). The major champion of the pragmatic approach to theoretical virtue is van Fraassen, who claims that 'the virtues claimed for a theory [beyond empirical adequacy and strength]...do not concern the relation between the theory and the world, but rather the use and usefulness of the theory: they provide reasons to prefer the theory independently of questions of truth' (1980: 88). The primary motivation for this view is that the virtues are functions of our interests and concerns – a theory that answers more questions about things that matter to us than its rivals will be preferred, but this preference is not guaranteed to be because the theory is more likely to be true than its rivals.

Kriegel (2013) argues that even if theoretical virtues might be truth conducive in the context of scientific theorising, there are particular reasons to think that concerns like van Fraassen's will apply in the metaphysical case. First, if virtues such as simplicity are to be truth conducive, an objective measure of simplicity is required. It is hard to see what such a measure could consist in (cf. Kriegel, 2013: 19). Presumably not in terms of explanatory simplicity (given the pragmatic nature of explanation) or in terms

of parsimony (because there is no particular reason to suppose that the world itself exhibits a kind of bareness: the preference for desert landscapes appears to be part of descriptive and not revisionary metaphysics).

Second, it seems unfounded (or even 'absurd', cf. van Fraassen, 1980: 90) to think that the world is more likely to be simple than it is to be complicated. In general, it seems far more likely that virtues such as simplicity are virtuous in so far as they make things easier to understand, or to explain, or able to express in fewer sentences, than that they are truth-conducive. Ease of understanding and communication are virtues that are relative to our current mental lives, and not necessarily to mind-independent reality. The onus is on the defender of the purported link between simplicity and truth-conduciveness to give a satisfying explanation for why simple theories are thus preferred, but not merely on grounds of pragmatism.

For van Fraassen, what it is to have a good explanation is largely just having a theory with virtues like simplicity to recommend it (1980: 94). The most virtuous theory of grounding is thus the one that we should accept, but not because it is any more likely to be true than its rivals. The point here is that Fine's appeal to theoretical virtue is not guaranteed to recommend a particular system of ground for the right sorts of reasons. If we prefer a particular system of ground on the basis of theoretical virtue, where those virtues are taken to be pragmatic, it follows that we should not find such preferences to provide compelling insights into the structure of reality with respect to grounding.

The remaining virtues mentioned by Fine are breadth, coherence and non-circularity. It is not clear from Fine's paper exactly what he means by breadth, but a common Lewisian approach takes something like breadth to be subsumed under theoretical strength, where strength is a matter of explanatory power. I assume that a theory is to be considered broader the more it is able to explain. Again, this will be subject to concerns that the explanatory power of the theory is a measure that is both not independent of context, and that does not necessarily render the theory more likely to be true than its rivals, which we might take to have less explanatory power (because, for example, the things the rival theory is able to explain are things we are less interested in).

Coherence or cohesion of a theory is a measure of entailment and probabilistic relations between beliefs in the system (or elements of the theory) (Kriegel, 2013: 23). The assumption is that more coherent theories are more likely to represent objective reality. But there are good reasons to reject that assumption. It is a common criticism of coherentists about justification that one might devise a perfectly coherent system that nevertheless ‘floats-free’ of truth, and thus that coherence is not in itself truth conducive. Sophisticated works of fiction can be constituted by extraordinarily coherent systems of statements, but are not linked to truth.

Kriegel (2013: 23) draws attention to recent proofs in Bayesian confirmation theory (Bovens and Hartmann, 2003, and Olsson, 2005) that he takes to underscore the sorts of observations made above. These purport to demonstrate that more coherent theories are not likelier to be true – coherence is no guarantee of correctness. (These proofs require certain assumptions to go through, notably that the degree of coherence of a system of beliefs is determined by the probabilities of the belief’s propositional contents, and that the individual probability of each belief’s propositional content is independent of that of any other’s. These assumptions seem *prima facie* plausible, though to defend them here would be beyond the scope of the thesis.)

It certainly does seem, however, that a coherent belief system is a benefit for the agent who holds it. There are practical advantages to a coherent system of beliefs in terms of justification, ease of drawing inferential connections, and lack of cognitive dissonance. Coherent belief systems are of pragmatically useful, but are not the path to discovering features of mind-independent reality. That a system of statements is coherent does not recommend the truth of those statements.

The upshot of all of this is that it is unclear how we ought to decide the answers to considerations about which explanations make for the best systems in the sense that Fine intends. We may be able to identify which are most pragmatically beneficial, but this will not satisfy the realist about grounding. Since the only way to evaluate grounding claims seems to be by appeal to intuition and explanation, neither of which are guaranteed to represent mind-independent reality, it would be better to take these appeals to intuition and explanation to justify making grounding claims without realist import if we can give a sensible antirealist story. As I will argue in the following sections, grounding claims as understood

in an antirealist spirit are the kinds of statements that appeals to intuition and explanatory considerations are both well-placed to justify. First, I wish to make one further comment about metaphysical explanation.

Recall that for Fine, grounding is the tightest explanatory connection between two propositions; ‘the ultimate form of explanation’ (2001: 16). I argued in Chapter 4 that this notion of metaphysical explanation is problematic, but I wish to claim here that its problematic nature can help to motivate antirealism about grounding. Mackie (1977) argues that to countenance objective moral properties requires us to posit things ‘of a very strange sort, utterly different from anything else in the universe’ (1977: 38). If there were such a thing, Mackie claims, we would further require a special faculty of perception to enable us to track it. The metaphysical queerness of moral properties appears, for Mackie, to consist in their motivational character. If moral judgements are intrinsically action-guiding, then the reasons for performing an action are independent of the agent’s desires.

The argument against explanatory considerations motivating realist grounding claims can perhaps be cast as something like a form of the argument from queerness. If there were such a thing as the ultimate, most metaphysically satisfying explanation of a given proposition (or phenomenon), then presumably that would be the explanation that we ought to accept. There is some normative force to a realist conception of explanation and in particular of *ultimate* explanation, and that normative force is queer in something like Mackie’s sense. The objectively correct explanation of some phenomena might come apart from the interests of the agent in seeking that explanation, and so the objectively correct explanation that ought to be accepted might seem to the agent to be unsatisfying.⁴⁸ This is analogous to the way in which, if moral properties have motivational force, an agent’s reasons might come apart from his desires.

We can bring out what seems strange about the idea that we might be expected to accept an explanation that comes apart from our interests using an example from van Fraassen, which he in turn adapts from Aristotle. A father asks his son why the outdoor light is on, to which the son responds by

⁴⁸ If explanation is understood in a more epistemic way this problem doesn’t arise; in that case, the interests of the agent are built in to what is the ‘correct’ (or best) explanation.

explaining that electricity is reaching the bulb because the switch completes the electric circuit that connects the bulb to the power source (van Fraassen, 1980: 131). The father feels the son is being impudent, because the answer he sought was something like ‘because we are expecting company’. If there is an objectively correct (metaphysical) explanation, then presumably it is the former, and the father ought to end his enquiry there.⁴⁹ We nevertheless are presumably justified in feeling that, in this case, the former ‘objectively correct’ explanation was not relevant. If there is no place for relevance in Fine’s account of metaphysical explanation (or in other accounts that ties the notion of ground to that of metaphysical explanation), then there is a queer obligation to end inquiry when presented with the ‘objectively correct’, metaphysical explanation, irrespective of one’s own interests.

As in the moral case, there is also an epistemic element to the queerness exhibited by objective relations of grounding (or objective grounding as expressed by a sentential operator). Not only would objective grounding be metaphysically queer in the sense that it brings with it an obligation to accept explanations irrespective of the agent’s interests, but, like moral properties, grounding would seem to require a special kind of faculty capable of tracking it. I have argued above that the faculty of intuition falls short of this task. Realists about grounding thus owe an account both of the explanatory force of local instances of grounding, and of our ability to know about those instances.

2.3. Antirealism or eliminativism?

One might worry that the arguments given so far against a realist approach to grounding appear only to motivate a kind of epistemological scepticism about the notion. If grounding is not the kind of thing we are able to find out about, we ought to stop paying attention to it. Rather than attempting to offer an antirealist account of grounding, the best thing to do is to eliminate grounding talk from metaphysics (and elsewhere in philosophy). There are two main arguments motivating antirealism about grounding over eliminativism. The first appeals to the use of grounding idioms outside of philosophy. The second highlights the usefulness of grounding talk, both inside and outside of philosophy. The specifics and

⁴⁹ This is something of a placeholder; an ultimate metaphysical explanation is more likely to be at the level of fundamental physics than at the more macroscopic level discussed in the main text.

force of the second argument can only be properly appreciated once the details of various antirealist positions have been articulated.

That there is a folk notion of grounding is evident from everyday uses of locutions such as ‘because’, where context makes it clear that the term is being used to express something like ‘in virtue of’. Recall that the ‘because’ of the Euthryphro dilemma is assumed to be a use that appeals to grounding (e.g. Evans, 2012). The question is whether it is the will of the gods that *grounds* the goodness of some action, or the other way around. Conversations about why a particular artwork is beautiful are about what beauty is grounded in. When parents explain why an action is wrong to their children, they present the child with grounding claims about the wrongness of the act. Grounding claims abound, and ought not to be ignored.

Further motivations for continuing to use the grounding idioms come from philosophical investigation. In Chapter 1 (section 2.2.1) I explained how Rosen (2010) uses the notion of grounding to develop some general principles, and explores their interaction with other metaphysical principles that are taken to be on firmer ground. He gives accounts of, amongst others, the grounding of universal facts, the grounding-reduction link, the entailment principle at work in grounding claims, the grounding of modal truths, and the determinate-determinable link. Rosen’s strategy is to see whether, when put to philosophical use, the grounding idiom breaks down and/or generates confusion or incoherence. He finds that this is not the case; philosophical theses can be framed in terms of grounding with relative ease, and thus it appears that continued grounding talk will be philosophically useful. If it is possible to state plausible principles about grounding and to develop arguments involving them, as well as to frame discussable questions involving grounding (as it appears we can) it is hard to see what else would be required to legitimise grounding talk (Rosen, 2010: 135).

Antirealism, in contrast to scepticism about grounding, provides a means to reconcile the everyday occurrence of folk grounding claims with the troubling epistemology of the realist’s notion of grounding. Various forms of antirealism (discussed in the following section) promise to provide alternative understandings of folk claims without incurring the realist’s commitments.

3. Types of antirealism

Having provided some general motivation for antirealism about grounding, in this section I explore some of the forms that antirealism about grounding might take. Broadly speaking, the antirealist about grounding needs to provide an account of grounding claims that explains how grounding claims can be appropriate in the relevant situations, even though those claims do not succeed in referring to any mind-independent notion of grounding. In this section, I give very brief overviews of the major strains of antirealism, and apply them to grounding discourse. (One of these – error theory – does not come out as antirealist on the characterisation of antirealism I gave above, because on this view statements of the discourse are systematically made false by features of the world. I include it below because it is important to the development of fictionalist views that do conform to my characterisation, because the relevant statements are made true or false by something other than what is suggested by the surface grammar of the statements of the discourse.)

The different versions of antirealism discussed below (subjectivism; non-cognitivism; error theory; fictionalism)⁵⁰ give different accounts of the nature of the content of the claims asserted and evaluated in the practice. I identify expressivism and fictionalism about grounding as the most promising approaches to antirealism about grounding, and develop these two positions in more detail in the following two sections.

3.1. Subjectivism

Subjectivism in a given domain of discourse is often introduced by appeal to metaphor; where the objectivist about some discourse likens our activity in that discourse to astronomy or geography, the subjectivist likens it to sculpture or imaginative writing (Dummett, 1978: xxv). The metaphor of joint-carving might also help capture what is at stake here (see Joyce, 2009 section 5); where the objectivist (or realist) envisages her task as one of discovering nature's joints, the subjectivist takes herself to impose a non-compulsory conceptual framework onto the relevant portion of reality, which is undifferentiated in

⁵⁰ I do not mean to suggest that this is anything like an exhaustive catalogue of positions that have been labelled 'antirealist' – far from it. Rather it is a brief run through what might be considered stereotypes of the most influential antirealist approaches.

and of itself.

Note that subjectivism differs from relativism in the following way. Relativism about a given domain of discourse holds that statements of the discourse contain an essential indexical element, but this need not be anything subjective. A statement such as 'John is short' is relative to a context ('John is short *for a pro basketball player*' might be true, whilst 'John is short *for a jockey*' might be false) but there is nothing mind-dependent or subjective about tallness (see Joyce, 2009).

Subjectivists about a given domain of discourse hold that statements of the discourse are truth-apt, and that the truth conditions for sentences in that domain are mind-dependent. For example, a moral subjectivist might equate rightness with the property of being approved of, either by themselves, or by some particular group, or by everybody (van Roojen, 2014a, section 1.3). The truth value of moral judgements thus depends on people's psychological states, rather than on any mind-independent moral facts. The upshot of this is that 'murder is wrong' might be 'true for me' if I genuinely disapprove of murder, whereas the same statement might be 'false for you' if you lack such disapproval.

Subjectivists about grounding are committed to holding that grounding claims have mind-dependent truth conditions. The problem with developing a subjectivist approach to grounding claims is a problem that is common to subjectivists in all areas: the distinctive subjectivist thesis is that the truth values of the relevant claims are mind-dependent, but there is much work to be done to spell out what mind-dependence amounts to. It is undoubtedly a useful placeholder, but is too coarse-grained to do serious work (see Joyce, 2009, section 5). One consideration against attempting to do the necessary work to pursue the subjectivist line is that the relativism that seems to fall out of a subjectivist approach to grounding would preclude enjoyment of many of the advantages of grounding talk that we might hope to maintain despite the adoption of some form of antirealism about grounding. If claims like ' x grounds y ' become in the mouth of the subjectivist ' x grounds y for me', then it is hard to make sense of disagreement about grounding.

Though there are more sophisticated versions of subjectivism, I take it that the attractive features of a subjectivist approach are adopted by the more sophisticated versions of antirealism I discuss below, and so I will not develop the subjectivist approach to grounding in any more detail here.

3.2. Non-cognitivism

Non-cognitivists in a given domain deny that there are any properties of the relevant kind (rather than holding that any such properties are mind-dependent). According to non-cognitivists, utterances of sentences in the domain in question typically express non-cognitive attitudes, as opposed to anything with substantial truth-conditions. Both non-cognitivists and subjectivists might agree that moral utterances conventionally express non-cognitive attitudes (a subjectivist might think that when somebody says something is right, she really says that she approves), but the key distinction between the two positions is that for the non-cognitivist, moral judgements don't express beliefs, and/or are not truth-apt (van Roojen, 2014a). Where the subjectivist thinks we express our dislike for something by saying that we dislike it, the non-cognitivist thinks we do so by something akin to booing and hissing. Non-cognitivism comes in a number of different varieties, mostly characterised by differences in explicating the semantic function of expressions in the relevant discourse, and the nature of the mental states expressed by those who utter sentences of the discourse.

For the purposes of applying the commitments of a non-cognitivist view to grounding talk, the most interesting variants of the view are expressivism and quasi-realism (with which expressivism is often combined). Expressivism is a thesis about what is expressed by the claims made in a given area of discourse. Rather than taking such claims to be descriptive, stating facts about the world, expressivists hold that they are expressions of attitudes, and thus have to do with the dispositions, judgements, sentiments and/or cognitive structures of speakers. Expressivism is most often defended as a metaethical stance (e.g. Blackburn, 1984; Gibbard, 1990), where moral claims might be taken to express pro- or con- attitudes such as delight or disgust, or other non-cognitive attitudes such as planning or intending (see Gibbard, 2003). There are also examples of expressivists in other areas, such as expressivism about modal discourse (e.g. Blackburn, 1993; Thomasson, 2007; Schnieder, 2010) and epistemic expressivism (see Chrisman, 2012). Barker (2007), Macarthur and Price (2007) and Price (2011) have defended versions of global expressivism. Expressivists about grounding hold that grounding claims do not express beliefs, and the positive expressivist proposal might be fleshed out in different ways.

The important point is that for the expressivist, the ‘direction of fit’ runs from the speakers of the discourse to the world, and not the other way around. That is, terms of the discourse in question ‘guild and stain’ reality ‘with the colours borrowed from internal sentiment’ (Hume, 1975: 294). Just as the ethical expressivist ‘projects’ his or her sentiments onto the world, the expressivist about grounding ‘projects’ features of his or her cognitive architecture onto the world. I’ll discuss two such approaches in the next section.⁵¹

The quasi-realist program is one of explaining the obvious realist trappings of the discourse in question, and of vindicating the legitimacy of the practice of making judgements in the relevant domain about which one is an antirealist. In the case of grounding, that project is to justify realist-seeming features of grounding talk in the absence of a commitment to realism about grounding. Blackburn (1993: 185) identifies two routes to quasi-realism: the first ‘fast-track’ route involves securing a notion of truth to regulate attitudinal discourse and justifying adherence to propositional form in the domain to meet the demands of such a notion. The ‘slow-track’ alternative demonstrates the applicability of realist-sounding talk to a domain about which we endorse antirealism in a more piecemeal fashion, demonstrating for a number of different speech-acts (e.g. assertions, inferences, interjections) how realist-sounding talk is generated.

One challenge the quasi-realist must meet is that of explaining how a truth predicate (for example) can be legitimately applied to what are ultimately just expressions of attitudes. Once this challenge is met, a further difficulty comes in properly distinguishing quasi-realism from genuine realism, given that realists and quasi-realists will assent to the very same sentences (see e.g. Dreier, 2004). (I’ll discuss this in detail in section 7). The precise way in which the truth-predicate is to be secured will depend on the details of the rest of the non-cognitivist proposal, and so I will not discuss it further here. At the least, the combination of an expressivist proposal for grounding claims with a quasi-realist program of securing the trappings of realist discourse offers a promising suggestion for future developments of antirealism about grounding, to be discussed in section 4 below.

⁵¹ If this projection metaphor means the view I am interested in is best described as a form of projectivism, so be it. I am not inclined to quibble about terminology here.

3.3. Error theory

Like the non-cognitivist, the error theorist about a given discourse denies that the entities putatively referred to by fragments of that discourse actually exist. In the moral case, for example, both the error theorist and the non-cognitivist deny the existence of moral properties. The difference is that the error theorist, and not the non-cognitivist, takes utterances of sentences of the relevant discourse to be assertions. Typical and literal utterances of the sentences in question are genuine assertions which express cognitive belief-like states, but the error theorist holds that those assertions are systematically false. An error theorist about grounding therefore holds that grounding claims such as ‘singleton sets are grounded in their sole members’ is false because grounding is not a feature of mind-independent reality. The error theorist might then be an eliminativist about grounding talk, or she might offer a story about why it is acceptable to continue to engage in grounding talk.

Eliminativism about grounding talk is unattractive because it cannot retain any of the seeming advantages of engaging in such discourse. Furthermore, it requires convicting the many philosophers and ordinary English speakers who do make grounding claims of systematic error. This is thought to be problematic because it is so uncharitable to ordinary speakers. A better strategy for the error theorist about grounding discourse to assume is to adopt a form of fictionalism about grounding, which dampens the assertive force of the problematic utterances.

3.4. Fictionalism

Fictionalists can be characterised in terms of their commitment to two theses, one ontological and the other metaphysical (Eklund, 2011). The ontological thesis is held in common with error theorists and non-cognitivists; that the entities characteristic of the discourse in question do not exist. The linguistic thesis is that utterances of sentences in the relevant discourse are not (or ought not to be) attempts at literal truth, but instead are aimed at (or ought to be aimed at) truth within some kind of useful fiction. Alternatively, those sentences might serve some kind of representational purpose, where what is represented does not have to do with the entities whose existence is at issue.

Fictionalists might hold that the relevant sentences are uttered under the scope of an implicit ‘within the fiction’ operator, or with an attitude of pretence or make-believe (as in Yablo, 1998; 2001; 2005). Either singular terms in the discourse purport to refer, but fail to in fact refer to anything outside of the fiction, or sentences in the discourse convey content that does not have to do with (i.e. is not *about*)⁵² what is purportedly represented by the surface form of those sentences.

Fictionalists about grounding treat grounding discourse as a fiction, and tell a story about the usefulness of grounding talk that accounts for the development of the grounding fiction. In section 5 I discuss two approaches to fictionalism about grounding, the first a version where grounding talk is strictly and literally false, and the second where grounding claims convey content that does not have to do with grounding. Before discussing fictionalism in more detail, in the next section I develop two accounts of expressivism about grounding.

4. Expressivism about grounding

Expressivism can be characterised as the conjunction of two theses, one negative and the other positive. The negative thesis states that the vocabulary in question is not ‘descriptive, not belief-expressing, not fact-stating, not truth-evaluable, or not cognitive’ (Price, 2011: 88). The positive thesis says that the vocabulary is expressing a non-cognitive attitude. Expressivists about grounding (whether or not they are also quasi-realists) owe an account of the attitude that is expressed by grounding claims. I’ll introduce two versions of expressivism about grounding, which I’ll call ‘the credence model’ and ‘the intelligibility model’. Both expressivist models dictate that a particular, complex, non-cognitive attitude is expressed with the use of grounding locutions. The ability to express these attitudes in a relatively simple way is part of the motivation to adopt an antirealist rather than an eliminativist approach to grounding talk. The models I discuss here are starting points for more comprehensive expressivist treatments of grounding claims. Though I present them separately, I am somewhat optimistic that a hybrid model (yet to be developed) might combine the best elements of each.

⁵² The relation of ‘aboutness’ is borne by meaningful items to whatever they are on, or of, or whatever they address or concern. Sentences are true or false because of how things are where the subject matter of the sentence is concerned, and the subject matter of the sentence is what the sentence is about (see Yablo, 2014).

Recall that there is some dispute over whether grounding claims ought to be expressed using a relational predicate (e.g. ‘grounds’) or an operator or sentential connective (e.g. ‘because’). It might seem as though expressivism about grounding is better suited to the connective view, because unlike the predicate view, the connective view seems not to carry any ontological commitment to relations (see Clark and Liggins, 2012: 816). In fact, I think expressivism is consistent with either approach to formulating grounding claims. Both predicate and sentential connective can be understood as an expression of a non-cognitive attitude. Just as a moral expressivist takes the predicate ‘is wrong’ not to carry ontological commitment to a property of wrongness, a grounding expressivist can take the relational predicate ‘grounds’ to be free of ontological commitment to a relation of grounding.

4.1. Credence model

According to the credence model, grounding claims reveal particular features of an agent’s epistemic situation. The basic idea behind the model is that grounding claims express a non-cognitive attitude which is a function of the agent’s credences. At present this is quite a general claim, but we might make it more precise in a number of ways. I’ll focus on one way it might be made more precise, which has to do with *resiliency* (see e.g. Skyrms, 1977).⁵³ We can appeal to Popper’s ‘paradox of ideal evidence’ (Popper, 2005: 425) to illustrate what is at stake. Suppose you are presented with a coin, and asked to assign a credence value to the possibility that it will come up heads when tossed. You are unsure whether or not the coin is fair, but if it is biased, you don’t know in which direction. The rational thing to do is therefore to assign a credence of 0.5 to the possibility that the coin will come up heads. Compare a situation where you toss the coin 100 times and it comes up heads 50 times and tails 50 times. It is then rational to assign a credence of 0.5 to the possibility that the coin will come up heads on the 101st toss. Your additional knowledge in the second case is simply not reflected in your degrees of belief in the outcomes of the two experiments. Resiliency is supposed to capture the difference between the two cases; the rational degree of belief you have in the outcome of the second experiment is *more resilient* than the rational degree of belief you have in the outcome of the first.

⁵³ Here’s an alternative: grounding claims express the agent’s commitment to keep credence in the grounding belief higher than credence in the grounded belief.

The credence model has it that grounding claims express a non-cognitive attitude of comparative willingness to update or to revise one belief over another with which the former belief is connected (in a sense to be explained below), in the light of new evidence. An agent's utterance of the sentence 'Socrates' singleton is grounded in Socrates himself' expresses the agent's comparative willingness to revise or to update beliefs about sets, rather than revise or update his beliefs about Socrates. In other words, it expresses an agent's inclination to treat his beliefs about Socrates as *more resilient* than his beliefs about Socrates' singleton.

Note that exactly how the credence model is to be cashed out will be affected by whether we take grounding to be a relation between entities of various ontological categories (as do e.g. Cameron 2008; Schaffer 2009; 2010a) or only a relation between facts or propositions (as do e.g. Audi, 2012a; 2012b; Fine, 2001; 2012a). When grounding obtains between objects (such as sets and their members) we can take the relevant beliefs (i.e. the beliefs towards which we have the non-cognitive attitude expressed in grounding claims) to be beliefs about the existence and nature of those objects. If we take grounding to obtain only between facts or propositions, then the relevant beliefs have the same content as those facts or propositions. For example, the claim 'Socrates' singleton is grounded in Socrates' expresses an attitude towards beliefs about the existence and nature of Socrates and about the existence and nature of Socrates' singleton. Attitudes towards the same beliefs are expressed by the claim 'the proposition <Socrates exists> grounds the proposition <Socrates' singleton exists>' (and again for 'Socrates' singleton exists because Socrates exists').

It is an important part of the credence model that some kind of connection between the relevant beliefs is part of the non-cognitive attitude that is expressed with a grounding claim, and so an account of that connection is owed. I envisage this connection as being the kind of relation that obtains between beliefs connected in a Quinean web of belief (see Quine, 1960: 12-13; Quine and Ullian, 1970). The basic idea is that the sentences we believe to be true form an holistic web of belief, connected through relations of implication. If a sentence is jointly implied by a set of sentences we believe to be true, then we are obliged to believe that sentence is also true, or to change our minds about one or more of the sentences

that collectively imply it (Quine and Ullian, 1970: 41). Grounding claims express our willingness to change our minds about one or other of these beliefs or sets of beliefs in parts of this holistic network.

Towards the centre of the web of beliefs are sentences that express grounding propositions, while sentences that express grounded propositions are towards the periphery. More evidence is required for us to revise the beliefs at the centre of the network – the grounding beliefs – than is required for the revision of beliefs on the periphery – the grounded beliefs, because the beliefs that are more central to the network are those that are more resilient. Because of the structure of the web, beliefs at the centre *explain* (in a sense that is agent-relative) the more peripheral beliefs, and so grounding is, in the appropriate sense, an explanatory relation. Our belief in the existence of Socrates is less open to revision than our belief in the existence of the set that has Socrates as its only member. Our belief that snow is white is less open to revision than our belief in the truth of the proposition <snow is white> (beliefs about propositions are less central than beliefs about snow). Note though that grounding claims do not express *propositions* such as <my belief in Socrates is less open to revision than my belief in the set that has Socrates as its only member>; this would be an expression of a cognitive state. The attitude expressed is a non-cognitive attitude of willingness to change one's mind about the grounded belief rather than about the grounding belief in the light of some new evidence.

Recall that theories about grounding and fundamentality can be seen as analogous to theories of justification in epistemology. Just as metaphysical foundationalists take entities (or perhaps just facts) to form well-founded grounding chains that terminate in the ontologically basic, epistemic foundationalists take justification to be inferred along linear chains from foundational beliefs for which no further justification is available. Just as infinitists about justification hold that justification is inferred among linear, non-terminating chains of belief, the metaphysical infinitist holds that entities are grounded in non-repeating, infinite chains. The metaphysical interdependence I argue for in Chapter 3 is analogous to coherentism in epistemology. Expressivism about grounding on the credence model can explain the apparent analogy between theories of fundamentality and theories of justification – grounding claims express attitudes towards beliefs, and the attitudes agents have towards their beliefs are closely related to

the ways those beliefs are justified. In general, the better justified a given belief, the more resilient that belief.

Expressivism about grounding on the credence model thus explains the intuitive pull of grounding claims (our intuitions about grounding reflect which of our beliefs are most resilient) and the explanatory nature of grounding claims. Successful explanations appeal to beliefs that exhibit high degrees of both resilience and justification. We should thus expect that grounding claims reveal a perceived explanatory connection between the relevant beliefs on behalf of the agent, such that the agent will be more willing to revise beliefs that are explained by more resilient beliefs than she will be to revise beliefs that explain less resilient beliefs.

Before moving on to consider the second version of expressivism about grounding, I will highlight and offer quick suggestions for responding to two worries for the credence model. The first concerns accounting for disagreement about grounding, and the second concerns some specific cases on which the credence model might seem to get the wrong result. A more general worry for expressivists – the Frege-Geach problem – is discussed in section 4.3.

One might worry that the credence model leaves no room for disagreement about grounding. Suppose Thomas says to Chrissie that an act's being morally obligatory is grounded in the will of the gods, and Chrissie responds that the will of the gods is grounded in an act's being morally obligatory. Though this looks on the surface as though there is disagreement about the truth value of some proposition (such as the proposition <an act's being morally obligatory is grounded in the will of the gods>), expressivism dictates that no proposition is expressed by an utterance of the relevant sentence. There is therefore no relevant proposition about the truth value of which Thomas and Chrissie can disagree.

According to the credence model, Thomas expresses greater willingness, in the light of new evidence, to revise or update his beliefs about what the gods will than to revise or update his beliefs about what is morally required, and Chrissie expresses greater willingness to revise or update her beliefs about what is morally required than to revise or update her beliefs about what the gods will. Thomas and Chrissie just have different attitudes towards their beliefs; there is no disagreement here. But it seems as though people *do* disagree over grounding (for example, when they argue about the Euthyphro dilemma) and so

if we are to preserve the intuition that there can be genuine disagreement about grounding, we need to provide an account of that disagreement.

Expressivists about normative discourse explain apparent disagreement as disagreement in attitude, spelled out in terms of clashing psychological states or conflicting advice, rather than in terms of truth conditions (Dreier, 2009: 83). The disagreement is thus like a disagreement about planning; disagreeing about what to do. A similar story can be told about disagreement about grounding. We can understand Thomas and Chrissie's disagreement as a disagreement about how to update or revise their beliefs in the light of new evidence. Thomas plans to revise his beliefs about what is morally required before revising his beliefs about what the gods will, and Chrissie plans instead to revise her beliefs about what the gods will before she revises her beliefs about what is morally required. We cannot explain Thomas and Chrissie's disagreement in terms of disagreement about the truth value of any proposition; instead it is a disagreement about what to do. (Compare the kind of disagreement a couple might have about where to go for dinner – one suggests an Italian restaurant, and the other responds 'no! Let's go for Lebanese food'. The couple disagree, and their disagreement is a disagreement in attitude.)

A second worry for the credence model is that there seem to be some grounding claims that it is not well-placed to account for. Proponents of grounding often take the claim 'mental states are grounded in brain states' to be a canonical example of grounding. Most speakers, however, would presumably report that they are more willing to update or revise their beliefs about mental states in the light of new evidence than to update or revise their beliefs about brain states. If the credence model is correct, the grounding claim expresses the opposite attitude. We are certainly more familiar with beliefs about mental states than with beliefs about brain states, but this difference in familiarity ought not to correspond to a difference in resilience. And on closer reflection, I don't think it does.

Under some circumstances (namely when we walk about other people's brain states and other people's mental states) I think it clear that beliefs about brain states *are* more resilient than beliefs about mental states. If we are looking at a scan of S's brain, and listening to her report that she is in pain, I think the belief formed on the basis of the scan ought to be more resilient than the belief formed on the basis of S's report. The methods we have for forming beliefs about mental states (testimony, and inference from

behaviour) are far more subject to error than the methods we have for forming beliefs about brain states (interpreting a scan, or believing on the basis of testimony from somebody else who has interpreted a scan). There are a number of other beliefs we might acquire that lead us to revise our belief that S is in pain (for example, if we have reason to doubt the reliability of her testimony, or we come to worry about the subjective character of pain, or about the way in which different people might report pain at different thresholds, or become concerned that S might mistakenly believe herself to be in pain). To revise our belief about S's brain state we would have to be presented with contradictory scans (or perhaps contradictory interpretations of the scan). People who make claims like 'mental states are grounded in brain states' are generally philosophers who are alert to these issues, and so we should not be surprised to find them expressing them using grounding locutions.

So what about the first person case? My pain is grounded in my being in brain state B, and I am surely less willing to update or revise my belief that I am in pain than my belief that I am in any given brain state. Here I think we can appeal to the idea that when making metaphysical claims such as the claim that mental states are grounded in brain states, the third person perspective has priority. Thus, the claim that my pain is grounded in my brain state is a straightforwardly follows from the claim that pains *in general* are grounded in brain states (or that mental states *in general* are grounded in brain states).

4.2. Intelligibility model

According to the intelligibility model, grounding claims reveal the comparative intelligibility (for an agent) of the relevant propositions, and any conceptual connections between them. As with the credence model, the intelligibility model of grounding claims can be adopted whilst ascribing to the view that the relata of the grounding relations are facts or propositions, or the view that they are entities of various ontological categories. On the latter view, it is the intelligibility of propositions that take the existence and nature of the relevant entity as their subject matter that is at issue. If grounding relates only propositions, then the intelligibility model straightforwardly concerns those propositions.

So, for example, an agent's utterance of the sentence 'Socrates' singleton is grounded in Socrates himself expresses the greater complication for the agent in grasping propositions like '<Socrates'

singleton exists> as compared with grasping propositions such as <Socrates exists>, and a conceptual connection between the propositions such that grasping the proposition that Socrates' singleton exists requires that the agent first grasp the proposition that Socrates exists. (Of course, the grounding claim does not express the belief (on the part of the agent) 'propositions like <Socrates exists> are easier to grasp than propositions such as <Socrates' singleton exists>', rather the grounding claim expresses the non-cognitive attitude concerning the relative intelligibility of the propositions to the agent).

Here's another example. Disjunctions are usually said to be grounded in their true disjuncts. In the mouth of the expressivist, an utterance of the sentence 'P∨Q is grounded in P' expresses the increased conceptual complication involved in grasping the disjunctive proposition over grasping the true disjunct, as well as expressing the agent's recognition that grasping the former involves grasping the latter.

Like the credence model, the intelligibility model can account for the presence of our deeply held intuitions about grounding. Because grounding claims are expressions of relations between our concepts and our attitudes towards the relative intelligibility of the propositions that involve them, we should expect that consulting our intuitions (which reveal features of our conceptual scheme) would generate apparent justification for our claims about grounding. Expressivism on the intelligibility model also offers an explanation for the way in which grounding claims seem to be explanatory. We say that snow's whiteness grounds the truth of the proposition <snow is white> because it is easier for us to understand that snow is white than it is for us to understand that the proposition <snow is white> is true. We feel that it is snow's whiteness that *explains* the truth of the proposition <snow is white>, because understanding the former helps us understand the latter.

As with the credence model, proponents of the intelligibility model owe an account of the nature of seeming disagreement about grounding. Here, the disagreement is fairly straightforward. Agents espousing different grounding commitments have different attitudes towards the relative intelligibility of the relevant propositions. This disagreement is similar in kind to the disagreement people might have over whether Brussels sprouts taste nicer than cabbage or not; it is a disagreement about attitudes, and not about propositions.

The intelligibility model must also offer an account of how to handle the claim that mental states are grounded in brain states. *Prima facie*, it seems that the intelligibility model would get this the wrong way around; we can grasp the proposition that S is in pain without any knowledge of the relevant neuroscience. It's not clear how the defender of the intelligibility view should respond to this objection, but here is one possible line of response. Many philosophers of mind who would make this sort of grounding claim think that the character of our experience is conveyed using 'phenomenal concepts'. Advocates of phenomenal concepts generally regard such concepts as exhibiting an unusually direct semantics, such that they are not *a priori* connected with other concepts. It would therefore be a mistake to think that a given phenomenal concept 'includes' the relevant physical concept with which it shares a referent. This is often taken to account for the 'explanatory gap' that appears to exist between (propositions about) mental states and (propositions about) physical states that are taken to ground them. Perhaps the friend of the intelligibility model can appeal to this explanatory gap to argue that it is to be expected that defenders of the relevant views espouse grounding claims that run in a different direction than we expect; when they make these claims, they make the claims that are entailed by their theory, rather than that accurately express their attitudes.

Underlying both of the models discussed above is the idea that grounding claims reveal something about the agent's relative levels of comfort with the relevant propositions, be this expressed in terms of the resiliency of their beliefs or the extent to which the agent finds those propositions intelligible. In both models, there is an explanatory connection between the two propositions such that the agent's attitude towards the grounded proposition is accounted for by their attitude towards the grounding proposition. As I mentioned above, the most promising expressivist account might be a hybrid that combines the best elements of both models discussed above, but this remains a project for future research.

4.3. The Frege-Geach problem

Familiar problems facing expressivists in other areas will also be faced by expressivists about grounding. Perhaps the best known example is the Frege-Geach problem (Geach, 1965: 463) which raises a

difficulty for giving an expressivist treatment of sentences when they are embedded in more complex sentences.⁵⁴ Consider the following example of a seemingly valid inference about grounding:

- (1) If singleton sets are grounded in their members, then {Socrates} is grounded in Socrates
- (2) Singleton sets are grounded in their members
- (3) Therefore, {Socrates} is grounded in Socrates

On an expressivist treatment, there is no proposition expressed by the claim ‘singleton sets are grounded in their members’, only an attitude. An expressivist about grounding will treat (2) as an expression of (for example) lesser willingness to update or revise beliefs about the existence of any given entity than to update or revise beliefs about the existence of the set which has that entity as its sole member. But it looks as though an agent who does not share such an attitude ought nevertheless to be able to assent to (1), and therefore that the embedded grounding clause in the antecedent of the conditional in (1) has a different content to the same clause as it appears unembedded in (2). We don’t ordinarily expect linguistic expressions to have different contents in embedded and unembedded contexts, and moreover, any difference in the contents of the expressions renders the seemingly valid inference invalid.

In the case of expressivism about moral discourse (where most work on expressivism has been conducted) the most influential responses to the embedding problem are as follows. First, one might adopt a minimalist theory of truth (see e.g. Horwich, 1993; Stoljar, 1993). For example, a minimalist might hold that collected instances of the schema: “‘S’ is true iff S” together imply everything there is to know about truth (Dreier, 2004: 26). Minimalism about related notions makes more realist-sounding discourse available to the antirealist. For example, for propositions: ‘S’ expresses the proposition that S. The expression ‘the proposition that’ is merely a device for forming noun phrases out of sentences so that generalisations are possible (see Dreier, 2004: 26, and section 7.1 below). Since grounding sentences are meaningful and can be embedded into that clauses, there are propositions about grounding.

⁵⁴ A detailed discussion of the embedding problem and responses that have been made to it would take us beyond the scope of this chapter. Here I only briefly mention the major styles of response.

In order to solve the embedding problem for grounding, it must be appropriate to call sentences about grounding true, for those sentences to embed grammatically in more complex sentences in the usual ways, and for us to have a relatively clear conception of the conditions on the appropriate use of sentences about grounding. The first requirement is secured by minimalism about truth, and the way in which we make grounding claims are evidence of the second. The positive expressivist proposals secure the third. We can explain the truth values of more complicated truth-functional embeddings by showing how their truth values are functions of the truth values of the component sentences. The minimalist expressivist can give minimal truth conditions for the component sentences and can thereby give an account of the meaning of the sentence whether asserted or embedded. (c.f. van Roojen, 2014b.)

The second available response to the embedding problem is to move to a hybrid expressivism that uses the ‘descriptive meaning’ of the relevant utterances to explain logical relations between sentences in the domain (see e.g. Ridge, 2006). Non-cognitivists are generally happy to grant some kind of secondary descriptive meaning to the sentences about which they are non-cognitivist (van Roojen, 2014b). For example, they might take the sentence ‘murder is wrong’ to express the speaker’s disapproval of murder (a non-cognitivist meaning) but also to predicate the property of moral wrongness of the act-type murder (a descriptive meaning). Hybrid expressivists give that descriptive meaning nearly equal status, and claim that (somehow or other – proposals differ, and do not concern us here) the descriptive component of meaning accounts for logical relations between the relevant sentences (see van Roojen, 2014b).

A third response exploits relations not between the contents of sentences, but between the sentences themselves (explaining the validity of inferences such as that introduced above). One influential version of this strategy for responding to the embedding problem was proposed by Blackburn (e.g. 1984) who posits a ‘logic of (higher-order) attitudes’ towards accepting certain attitudes. Applied to our inference about grounding, we might claim that an utterance of the statement in (1) expresses the agent’s attitude towards her willingness in the light of new evidence to update or revise beliefs about Socrates’ singleton over updating or revising beliefs about Socrates – she takes it that that attitude is implied by her willingness in general to, in the light of new evidence, update or revise beliefs about singleton sets over updating or revising beliefs about their sole members. The agent’s utterance in (2) expresses her greater

willingness in the light of new evidence to revise or update her beliefs about the existence of singleton sets than to revise or update her beliefs about the existence of their sole members. Given the resilience of her beliefs as expressed in (1) and (2), it would be irrational for the agent not to hold the attitude expressed by (3); greater willingness in the light of new evidence to revise or update her beliefs about the existence of {Socrates} than to revise or update her beliefs about the existence of Socrates.

These responses to the embedding problem work by complicating the semantics of grounding claims so that they can sustain the problematic inferential connections. Minimalism and/or positing a logic of attitudes allows the sorts of inferences that are constitutive of realist discourse to go through. Blackburn's quasi-realist program is the general project of securing the ability for the antirealist to legitimately engage in realist-sounding discourse. I will have more to say about quasi-realism and its consequences in section 7.

One final possible solution to the embedding problem I will mention is to distinguish between local expressivism – expressivism about a given area of discourse, and global expressivism – an account not restricted to a particular area of discourse, and to adopt the latter. The recent champion of global expressivism is Huw Price (see also Barker, 2007, and the Sellarsian version of expressivism advocated by Brandom e.g. 2000). Price argues that 'whatever story the quasi-realist tells about the genealogy and functions of the "factual" character of (say) moral language, the same story is likely to work for other cases, too' (Price, 2011: 97). The idea is that there is pressure to give a uniform account of declarative speech acts, and that that entire story can be non-representational. At the local level, there is diversity between vocabularies – Price wants to maintain that (for example) the states projected in association with moral concepts are different from those projected in association with modal concepts and that they play different roles in people's lives (2011: 98). There is a distinction between domains of discourse, but the distinction is not between representational and non-representational vocabularies, because all vocabulary is non-representational.

Global expressivism seems to avoid embedding problems because it removes a distinction between parts of language that are to be given a cognitivist treatment (i.e. are to be considered truth-apt) and those that are not. There is of course a worry that global expressivism throws the baby out with the

bathwater – in attempting to avoid troubling metaphysical commitments in the case of one area of discourse (such as morality or grounding) we are led to adopt an antirealist position with regards to everything. This difficulty is discussed by Price (2011), who argues that global expressivists can accept the idea that some vocabularies are more in the business of ‘tracking the world’ than others, but that this limited bifurcation is not to be made at the semantic level. Global expressivism is an intriguing but radical position, and since I think there are alternatives to global expressivism open to the expressivist about grounding I will not discuss it further here.

5. Fictionalism

Fictionalism about a given domain of discourse is (roughly) the view that claims within the relevant discourse are best interpreted as fictional: as either accompanied by an attitude of pretence or make-believe, or as bound by an implicit ‘within the fiction’ operator. Different fictionalist approaches treat the sentences in the domain of discourse about which they are fictionalist in different ways. Here I’ll introduce two types of fictionalism about grounding. The first is a Field-style fictionalism whereby sentences in the relevant domain contain genuine referring expressions that purport to refer to the entities under dispute. Since there are in fact no such entities, those terms fail to refer and so sentences in the domain are strictly and literally false. The fictionalist then tells a story about how such sentences are nevertheless assertible. The second is a Yablo-style fictionalism which has it that sentences in the domain are true, but are to be understood metaphorically rather than literally.

Cutting across this distinction between styles of fictionalism is a distinction between hermeneutic fictionalism and revolutionary fictionalism (see e.g. Stanley, 2001: 36; Eklund, 2005: 557). Hermeneutic fictionalism is a thesis about the actual nature of the discourse – it holds that statements made within the discourse do not aim at the literal truth but only appear to pretend to do so; normal use of the discourse involves pretence. When I say ‘Harry Potter’s best friends are Hermione Granger and Ron Weasley’, I do not mean to commit myself to the genuine existence of any of those people. I intend my statement to be understood within the fiction of Harry Potter (though this intention need not be explicit, or even conscious). By contrast, revolutionary fictionalists admit that the relevant discourse does in fact involve

literal reference to the problematic entities, but hold that we *ought* to use the discourse in such a way that the reference is only within the pretence. When engaging in the relevant discourse we ought only to make pretend-assertions, and the point of engaging in the discourse would be achieved if we made only pretend-assertions (see Eklund, 2011, section 2.2).

5.1. Field-style fictionalism about grounding

If the positive, atomic sentences of a domain of discourse are systematically false (e.g. because they make reference to non-existent entities) then we might initially assume that we ought to discourage people from using those sentences. If, however, we can provide an account of why those sentences, though false, are useful and assertible, we need not seek such a change. Fictionalists inspired by Field (e.g. 1980) take up the task of providing an account of the assertability of such sentences in the face of eliminativist antirealist arguments in a given domain (i.e. arguments that show that sentences in the domain are strictly and literally false). Field (1980) goes to great lengths to show that numbers are in fact dispensable to science. Strictly speaking, science doesn't need numbers, but it is far easier to do science with them than it is to do it without them. Number-talk is a very useful way to simplify science, and so it is worth engaging in even if there are in fact no numbers. A Field-style fictionalist about grounding must show that even if grounding talk is strictly and literally false because there is no mind-independent relation of grounding, it can still be worthwhile to make grounding claims.

Before I discuss how the fictionalist about grounding might carry out this task, we should note that if 'grounds' refers to 'grounds in the fiction' (and so fictionalism is a semantic thesis – a thesis about the referents of grounding locutions) then it might seem as though fictionalists about grounding ought to prefer the predicate over the operational formulation of grounding claims so as to give a fictionalist treatment of the predicate 'grounds' whilst leaving the other elements of the sentence untouched. I don't think this is required by the Field-style model, however. The Field-style fictionalist about grounding can take the 'because' in an operational formulation of grounding claims (e.g. $P \vee Q$ because P) to pick out something like 'because_{fiction}' – a fictional sentential operator (which might operate on non-fictional sentences). (Compare 'it-is-not-the-case-that_{HarryPotter} David Cameron is prime minister'.) It is also

consistent with Field-style fictionalism to take entire sentences to mean something bound by an implicit ‘within the fiction’ operator. So, an utterance of the sentence ‘Socrates’ singleton is grounded in Socrates himself’ means ‘within the fiction of grounding, Socrates’ singleton is grounded in Socrates himself’, or an utterance of ‘ $P \vee Q$ because P ’ means ‘within the fiction of grounding, $P \vee Q$ because P ’.

The most plausible Field-style fictionalism about ground is a revolutionary fictionalism. It is somewhat implausible to suppose that philosophers engaged in grounding talk have been aiming at truth within the fiction of grounding all this time, but none of them have ever mentioned it. It is also hard to make sense of the realist rhetoric of influential authors such as Fine and Rosen if hermeneutic fictionalism is a proper account of the practice.⁵⁵

The revolutionary fictionalist owes an account of how the benefits of engaging in grounding discourse can be reaped in the face of the literal falsity of claims about grounding. The first step is to rehearse the benefits of engaging in grounding talk in the first place. As we have seen (e.g. Chapter 1, section 2.2, and section 2.3 above) proponents of grounding such as Rosen (2010) and Schaffer (2009) argue that much is gained by incorporating the idioms of grounding into our ‘analytic toolkit’ (Rosen, 2010: 110). Doing so enables us to restate and sometimes to offer a solution to familiar puzzles (I discuss an example in section 7, and demonstrate that the solution goes through even if we are antirealist about grounding), and to systemize and simplify discourse in numerous areas of philosophy.

Even though grounding claims are literally false, it can be pragmatically useful to behave as though they are true. If we can give a unified antirealist treatment of grounding claims then we can still appeal to grounding to systematise and simplify. In Field’s mathematical fictionalism, the appeal of number-talk is the role it plays in simplifying scientific discourse. There is no need to be realist about number-talk for us to take the discourse to play a simplifying role. Similarly for grounding talk; if grounding talk really does

⁵⁵ Perhaps this is a little quick. It can be a part of hermanautic fictionalism that speakers don’t really understand the nature of what they are saying, and so the fact that proponents of grounding talk don’t take themselves to be making pretend-assertions about grounding ought not to be considered a decisive objection. Nevertheless, I think this argument has more weight as applied to philosophers working on grounding than it does in the moral case where we might think it more plausible that ordinary speakers could be somewhat in the dark about the nature of their moral utterances. As both hermeneutic and revolutionary fictionalists can give similar motivations for engaging in the pretence, I’ll assume that most of what I say below applies to both accounts.

have a role in simplifying and systematising metaphysics, then this is an advantage that even a fictionalist about grounding can enjoy.⁵⁶

Here is one kind of example. Recourse to grounding talk in a given domain of discourse provides a simplification that is often welcome, especially to those unfamiliar with the relevant theories. When the moral theorist claims that moral facts are grounded in natural facts, she shies away from giving a thorough specification of how the moral facts relate to the natural facts. What she does do by using the grounding locution is convey something important about the nature of the connection – she tells us that there is one, that the connection is somehow explanatory, the direction the connection runs in, and that it matters in theorising about her discipline. All of this is possible in the absence of any mind-independent grounding relation.

Field-style fictionalism about grounding has a ready account of why it is that intuitions seem able to reveal facts about grounding. The fiction of grounding is created and sustained by those who appeal to it. In the technical, philosophical context, it has been developed and formalised by philosophers educated within the same tradition and engaging with one another's work. As co-authors of the fiction, we should expect those individuals making grounding claims to be able to appeal to their own intuitions as a fairly reliable guide to the truth of some relevant fictional claim. Disagreement about some claim within the fiction arises because the grounding fiction is a collective work, and so some authors have slightly different ideas about particular features of the fiction. (Imagine for comparison a book written by a group of people, where the authors disagree about the motivation for some action performed by the main protagonist.) Authors of the fiction are in a position to know the most salient explanation of any given fact within the fiction, and so appeals to explanation to make grounding claims are, like appeals to intuition, likely to be truth-within-the-fiction-conducive.

Appeals to collective authorship might be plausible in the case of the technical notion of grounding discussed in the philosophy room, but what of the everyday use of 'in virtue of' and cognates in folk

⁵⁶ Field (1980) works hard to show that mathematical theories are conservative over nominalistic ones. A fictionalist about grounding of the kind described above should ideally be able to demonstrate that anything we can say using grounding talk can also be said without it. The difficulty in accomplishing this task might be a reason to favour Yablo-style fictionalism about grounding. (Then again, perhaps the burden of proof is on the non-fictionalist, who must find some grounding claim that cannot be expressed in other terms.)

discourse? I think it plausible to suppose that this too is an instance of collective authorship. Any folk appeals to such terms ought to be (and to be interpreted as being) made in a fictional spirit, and the claims that involve those fictional terms are made on the basis of appeal to intuition and explanation, and are mostly veridical because they are made by the authors of the fiction. Thus, when a mother says to her child 'it's wrong to hit your brother because it hurts him', she makes the claim that the wrongness is grounded in the natural property of the pain caused to the child's brother, but her claim is only to be understood as true within the fiction of grounding. In making that claim, she contributes to the fiction.

The fictionalist about grounding might here appeal to the way in which intentions play a role in semantics in order to give an account of truth within the fiction of grounding. At least part of the way in which the referents of at least most expressions are determined is by the way in which the expression in question is used by members of the linguistic community. Analogously, truth within a fiction is determined (at least in part) by the intentions of the authors of the fiction. If enough people making grounding claims take moral properties to be grounded in natural properties, then it becomes true within the fiction of grounding that moral properties are grounded in natural properties (even if, in fact, there are objective moral properties).

This idea can be supported by the kind of social externalism found in Tyler Burge (e.g. 1986). Burge argues that the contents of some thoughts and beliefs are determined by the linguistic community of the speaker. The idea can be explained using one of Burge's arguments for his position, which runs as follows. Imagine an agent, Art, who goes to the doctor complaining of arthritis in his thigh. He is unaware that arthritis is in fact a disease of the joints, and so he expresses a false belief when he sincerely utters 'I have arthritis in my thigh'. We can then imagine a counterfactual scenario where Art has the same history and the same internal states, but grew up in a community where the word 'arthritis' is used to apply to a different disease, call it tharthritus, with includes rheumatoid conditions of the thigh as well as the joints. Now when Art says 'I have arthritis in my thigh' he expresses a true belief about tharthritus. Since neither he nor anybody else in his linguistic community has a concept of arthritis, it cannot be that his belief is really about arthritis. But because Art's internal states and history are identical in both cases,

what makes his belief true or false is dependent on the concept the linguistic community takes the word ‘arthritis’ to express.

We might be sceptical about the reach of Burge’s social externalism, but the area in which it has been taken to be most plausible is where the concepts involved are deferential (see e.g. Loar, 1990). Appeal to social externalism will therefore be particularly well motivated for defenders of grounding who are uncomfortable with the idea that folk grounding claims add to the fiction of grounding, on the basis that grounding is a semi-technical notion. If we take grounding to be a deferential concept and give a social externalist analysis of such concepts, we can explain how it is that grounding claims can be made in ordinary discourse, without such claims having much effect on the fiction of grounding itself. The folk can participate in the grounding story without themselves being very productive authors of it.

5.2. Yablo-style fictionalism about grounding

A different kind of fictionalism championed by Yablo (e.g. 1998, 2001, 2005) holds that assertions of sentences in the discourse about which one wishes to be fictionalist are analogous to ordinary instances of figurative speech. We use make-believe for serious descriptive purposes. To make as if to assert P (where P is some sentence in the domain in question) is to assert Q, where Q says that things are a certain way, and it is the way things are that makes it appropriate to assert Q. The real content of the asserted sentence P is in fact Q. Here is an example which Yablo attributes to Walton (1993: 40). We can convey information about the location of the Italian town of Crotone by saying it is on the arc of the Italian boot. We thereby assert the location of the Italian town, by making as if to assert that the town is on the arc of a boot.

On a Yablo-style model, sentences about grounding are representational aids, or myths, apt for conveying content that does not have to do with grounding. So what is this conveyed content? Again, there is room for divergent views, and answers to this question will be similar to answers to the question of which attitude is expressed on an expressivist view. (Though, of course, on the expressivist view a non-cognitive attitude is *expressed*, and on the fictionalist view some related content is *asserted*.) One quite attractive account of the conveyed content of grounding claims has it that grounding claims convey

information about comparative degrees of confirmation or support enjoyed by different propositions in a theory, and the entailment relations between those propositions (this is similar to the intelligibility model described above). On this understanding, grounding talk is a representational aid apt for describing features of theories, understood as complex, holistic structures of propositions. Propositions more central to the network are both better supported than, and able to support, propositions on the periphery. These more central propositions will be the grounding propositions (or will have as their subject matters grounding entities), and the grounded propositions (or their subject matters) are those that are less well confirmed or supported. Note that this sort of understanding of fictionalism fits very well with the metaphysical interdependence introduced in Chapter 3, which has it that entities can enter into mutually supportive relations of grounding. Another sort of account (closely related to the credence model described above) has it that grounding claims convey the comparative resiliency of different beliefs. Beliefs about grounding entities are more resilient than beliefs about grounding entities, and those beliefs are connected in a web of belief as described above. As with the expressivist models above, some hybrid of the two positions might make for an attractive account of the conveyed content.

Since Yablo-style fictionalism operates at the level of sentences, we might think that an approach to formulating grounding claims that treats grounding as a sentential operator is the best fit with this kind of antirealist view. In fact, I think the view is also well suited to deal with a view that formulates grounding claims using a predicate 'grounds'. A Yablo-style fictionalist about grounding can treat relational predicates like 'grounds' as representational aids in much the same way as they can treat mathematical language (and mathematical relations such as 'is the sum of') as representational aids. There is no particular problem for Yablo-style fictionalism about grounding in accounting for grounding claims formulated using a predicate.

If it is right that grounding talk is a kind of myth or a game, apt for representing content not to do with grounding, then the fictionalist owes an account of the development of this myth (cf. Yablo 2005). There are at least two aspects to this. The first concerns why the myth would develop at all, and the second concerns why this particular myth (with its seemingly distinctive correctness conditions) would be apt for conveying the appropriate content.

My story of the development of the myth of grounding hangs on the claim that the content conveyed by grounding claims is content that ordinary speakers have cause to convey. It is to be expected that such speakers might sometimes have a need to represent their attitudes to various propositions or beliefs and the connections between them, and it is evident from my discussion above that information about networks of propositions or about the resiliency of different propositions and the connections between them is complex, clunky, and hard to express. Given this, it is unsurprising that some kind of representational mechanism would develop as part of the language such that these attitudes towards propositions and beliefs might be expressed.

The next stage in the development of the myth of grounding is to grant that the use of ‘because’ and similar locutions by ordinary speakers does sometimes express the relevant kinds of attitudes. We have already encountered some examples. Claims like ‘you shouldn’t hit your brother because you will hurt him’ and ‘doesn’t the use of those colours make the painting beautiful?’ are apt for conveying the following kinds of content. In the first case, (for example) that the belief that hitting somebody will hurt them is more resilient than the belief that it is wrong to do so (and that the beliefs are connected in a web via explanatory implication relations as described in section 4.1). In the second, (for example) that propositions about colour are better confirmed than, and themselves support, propositions about beauty.

At this point we can explain the development of the theoretical notion of grounding as discussed by philosophers. Both accounts of the content conveyed by speakers making grounding claims are in some minimal sense explanatory, because they each involve implication relations between beliefs, or entailment relations between propositions. When an agent takes one proposition to entail another, or one belief to imply another belief, that agent takes the implied belief or entailed proposition to be *explained by* the belief or proposition that it is implied by, or that entails it. Explanatory notions are murky (see Chapter 4), and so my story of the development of the grounding myth has it that philosophers picked up on the complex and varied nature of terms like ‘because’ as used by ordinary speakers. In an attempt to systematise different explanatory notions, philosophers labelled this particular type of explanatory connection ‘metaphysical explanation’, and began to theorise about it. (The mistake that these philosophers made, as I see it, was misinterpreting the way in which ordinary speakers conveyed

information about beliefs or propositions and the connections between them for speakers making a claim about some more ‘objective’ form of explanation. This mistake is responsible for the difficulties in properly articulating a notion of metaphysical explanation (see Chapter 4.)

Philosophers attempted to sharpen and to codify this notion of metaphysical explanation in their talk of grounding. Just as mathematicians require an autonomous description of their imagined numbers (Yablo, 2005: 99), philosophers sought to discuss grounding as an independent notion. Both ways in which we might develop a Yablo-style fictionalism about grounding (by appealing to resiliency or to levels of confirmation) turn the grounding discourse into a kind of ‘game’, whereby participants in the game can come to treat the game itself as something that can become an object of study in its own right (cf. Yablo, 2005: 100). In much the same way as we can say things about the Italian boot (e.g. that it is high-heeled) and about numbers (e.g. that the sum of any two odd integers is always even), we can also say things about ground (e.g. that it is transitive). Interest in grounding itself is both about the game, and about the features of the world for which grounding is being used as a representational aid. So, the logic of ground conveys both something about the world and something about grounding itself. Though our ultimate interest is in describing and representing the world, we have a significant secondary interest in developing, revising and consolidating the ‘games’ we use to do so.

Philosophers thus themselves took over and elaborated a game of make-believe apt for conveying the kind of content conveyed in folk utterances of purported grounding claims. The philosophers’ game, having its roots in the folk practice, is correspondingly apt for conveying the same sorts of content. The make-believe of grounding makes it easier to systematise the contents of sentences about grounding, to reason about them, and to draw connections between them and other worldly facts (these points are made about mathematical fictionalism in Yablo 2005: 98). Without recourse to grounding talk, a speaker might well lack the expressive resources to convey the relevant information. She might be unaware that she is even in possession of the information, but still able to convey it by engaging in grounding talk (compare a child who is able to add 2 and 2 to make 4, but unable to understand that this might represent a general worldly fact like that whenever we have the same number of Fs as we do Gs, together

we have twice as many (FVG)s as we do Fs). Grounding talk thus plays an important role both in folk discourse and in philosophical theorising.

The second thing the fictionalist owes an account of is why it was the grounding myth (rather than some other myth) that developed for representing the relevant content. Some formal and structural features of grounding seem to directly reflect features of the content grounding talk conveys. If we take *a* to be a belief we think highly credence worthy, *b* to be inferred from *a* and to be worthy of slightly lower credence, and *c* to be inferred from *b* and worthy of slightly lower credence again, then we should be expected to be able to infer *c* from *a*, and to take *a* to be more credence-worthy than *c*. This accounts for the transitivity of grounding in the game of make-believe.

Beyond the straightforward reflection in grounding talk of features of the conveyed content, there may simply be no satisfying answer as to why this particular myth. We needed some way to represent the complex content that grounding talk conveys, and this way is as good as any other (and presumably better than some).

6. Comparative remarks

In the previous two sections I suggested two versions of expressivism about grounding, and two versions of fictionalism about grounding (where the second fictionalist model again bifurcated into two distinct accounts of grounding claims). My aim has not been to argue for any particular position (my interest is in the general possibility of antirealism about grounding) but rather to map out some forms of antirealism about grounding that might be candidates for further attention. In this section I make some brief comparative comments.

Both expressivism and fictionalism offer accounts of grounding claims that make such claims assertible and legitimate, even though both accounts also hold that grounding claims fail to pick out a notion of grounding ‘out there’ in objective reality. For the expressivist, the failure to correspond with reality, and thus the departure from realism, occurs at the level of content. Whilst the realist takes the content of grounding claims to have to do with real grounding relations, the expressivist thinks that the content of

grounding claims is some non-cognitive attitude of the agent making the claim towards a certain function of her credences, or towards the intelligibility of certain propositions, and the connection between them. Yablo-style fictionalism also operates at the level of content, such that what is conveyed by a grounding claim is content not to do with grounding. These non-representational notions of ground demand a non-factualist semantics. For the Field-style fictionalist, the departure from realism occurs at a different point. Both the Field-style fictionalist and the realist agree that the content of grounding claims genuinely represents grounding relations as obtaining, but the Field-style fictionalist holds that since there is in fact no such relation – grounding claims are best understood as fictional. There is a failure of correspondence with mind-independent reality.

It is possible that a Yablo-style fictionalism about grounding might also be combined with a form of expressivism about grounding in the folk discourse. This view would have it that the folk make grounding claims in order to express particular attitudes (towards beliefs or propositions) and that philosophers invented a game of make-believe apt for systematising and theorising about the conveyed content and about the game itself. Ordinary speakers, on this view, are not themselves participants in the game. The advantage of this kind of hybrid view is that it explains the sense in which grounding seems to be a semi-technical notion, used in more technical ways by philosophers than by ordinary speakers.

7. Realism, sophisticated antirealism, and grounding

Let's take stock. In this chapter I have mentioned some motivations for antirealism (section 2), outlined some antirealist accounts (section 3) and developed the most plausible of these as applied to grounding in some detail (sections 4 and 5). With these antirealist accounts in hand, in this section I want to discuss in detail one way that grounding talk has been put to theoretical use – to distinguish realism from sophisticated antirealism – and to argue that an antirealist approach to grounding is up to the task. I'll introduce the problem, and explain Kit Fine's solution. Given the antirealism about grounding outlined above, the original problem seems to recur – how are we to distinguish realist from sophisticated antirealist *about grounding*. I claim that the grounding theorist can answer this question, and so marking a

distinction between realism and sophisticated antirealism is an important theoretical role that both realist and antirealist in any domain can appeal to grounding talk to fulfil.

7.1. The problem of ‘creeping minimalism’

Fine (2001) invokes his notion of ground in order to solve a problem – a problem that Dreier (2004) calls the problem of ‘creeping minimalism’. Fine identifies a discontinuity between ordinary discourse, and the commitments we make in the philosophy room. An antirealist about mathematics might be willing to assert sentences such as ‘3 is a prime number’ (when helping her child with his maths homework, for example), despite denying that there are any numbers. Ethical expressivists assert ‘torturing children is wrong’ despite their maintaining that there aren’t any moral properties, and mereological nihilists are happy to say ‘this table is bigger than that one’ even though they don’t think there are any composite objects. I’ll call the nominalist, the expressivist and the nihilist antirealists about their respective domains of discourse.

The reason for the discontinuity between ordinary discourse and philosophical commitment is that antirealists in various domains have come to adopt the following two theses: *accommodation*, and *minimalism* (Dreier, 2004: 26). It would be a cost for nominalism about mathematics if the nominalist were to find herself unable to help her child with his homework in the light of her rejection of realism about numbers. Better that she find a way to accommodate number talk as it is used by ordinary speakers of her language, and that she do so in a way that is consistent with her nominalism. Accommodation is the process of taking on board this lesson, and making sure that, so far as possible, one’s theory is not in conflict with actual linguistic practice.

We encountered minimalism above (section 4.3); adopting minimalist accounts is one way expressivists have responded to the embedding problem. Recall that the thrust of the minimalist idea is as follows (Dreier, 2004: 26). There is nothing more to truth than collected instances of the schema:

‘S’ is true iff S

Similarly for propositions

‘S’ expresses the proposition that S

And for properties

x has the property of being F iff x is F

Minimalism ‘sucks the substance out of heavy-duty metaphysical concepts’ (Dreier, 2004: 26). Minimalist theses capture the thought that there is nothing more to saying that it is a fact that murder is wrong than emphatically saying that murder is wrong (Dreier, 2004: 25). Antirealists who adopt minimalism are happy to claim not just that torturing children is wrong, but also that it is true that torturing children is wrong, that they believe that torturing children is wrong, that it is a fact that torturing children is wrong, and so on. I’ll call antirealists who adopt both accommodation and minimalism ‘sophisticated’ antirealists. In Fine’s words, the propositions of the sophisticated antirealist are like zombies; just as philosophical zombies are creatures that look, talk, and behave exactly like us – displaying all the outward signs of consciousness but with no conscious experience – the sophisticated antirealist’s propositions exhibit all the usual trappings of realism but none of the substance (Fine, 2001: 12).

The problem of creeping minimalism is that in enabling sophisticated antirealists in various domains of discourse to recapture realist language (i.e. to achieve *accommodation*), minimalism threatens to make sophisticated antirealism indistinguishable from realism (Dreier, 2004: 26). There is nothing the realist can say that the sophisticated antirealist won’t themselves be able to agree with. We cannot distinguish between realist and sophisticated antirealist on the basis of the claims they are willing to endorse about the entities in question.

Here’s another way we might try to make the distinction. Quine (1948) tells us how to discover the ontological commitments of a theory. We begin by writing out the sentences of the theory in their canonical notation, in order to represent the underlying quantificational structure behind the surface grammar of the sentence. The canonical notation makes explicit which quantifiers occur in the sentence, and what their scope is. One is ontologically committed to all and only those things that are the value of a bound variable occurring in one or more of the sentences of the theory, and the set of sentences implied by the theory (Quine, 1948: 31-2). In essence, a theory is committed to whatever must exist for that theory to be true (Hookway, 1988: 25).

Because realist and sophisticated antirealist think that different entities must exist for their theory to be true, we might hope to use the Quinean criterion of ontological commitment to distinguish realist from sophisticated antirealist. In fact, this won't work. The sophisticated antirealist is willing to assent to sentences with the very same canonical notation as the realist, and so the Quinean criterion will be blind to the differences between them. Here's an example. In the mouth of both the realist and the sophisticated antirealist, the sentence 'this table is bigger than that one' is given the following Quinean paraphrase:

There is some x which is a table and some y which is a table and x is bigger than y

Or:

$\exists x \exists y ((x \text{ is a table} \ \& \ y \text{ is a table}) \ \& \ x \text{ is bigger than } y)$

The sophisticated antirealist hopes to assert (believe, embed, etc.) this sentence without accepting tables into her ontology. The ontological commitments of sophisticated antirealists come apart from the language of their best theory.⁵⁷

Of course, it is open to the Quinean who fails to recognise a difference in ontological commitment between the realist and the sophisticated antirealist on the basis of the sentences they are willing to assert to deny that there is really any such difference. Similarly for anybody who maintains that there is nothing more to realism than believing, embedding and asserting certain sentences one takes to be true. But this bullet-biting response is unsatisfying. Sophisticated versions of antirealism certainly originated as positions distinct from realism, denying what the realist affirmed. It was only via the processes of accommodation and minimalism that antirealists came to assert realist-sounding sentences. In doing so the sophisticated antirealists did not, by their lights, give up on anything characteristic of an antirealist position. I'll assume here that the relevant forms of sophisticated antirealism are indeed to be distinguished from realism, and our project is to work out how.

⁵⁷ One might think this is a problem for the antirealist rather than for the Quinean criterion. I'm not interested here in arguing over who the problem belongs to (and I suspect any such argument would beg questions on both sides). Instead, I'm interested in how a different method might more successfully distinguish realist from sophisticated antirealist.

Note that the problem of creeping minimalism belongs to neither the realist nor the antirealist, but to metaontology (see Dreier, 2004: 31). It is a problem for anybody who believes that realism and sophisticated antirealism about a given domain of discourse are distinctive positions, and that realist and antirealist espouse different ontological commitments. The project is to discover a way to discern those differing commitments.

7.2. Fine's 'Question of Realism'

Fine argues that strategies that have been proposed for distinguishing between realism and sophisticated antirealism are each subject to counterexample (see Fine, 2001 section 2). Most of the discussion is focussed on Dummett's proposal that realism for a given domain of discourse is a matter of conforming to the principle of bivalence. The main problem is that the proposal does not answer to our needs (Fine, 2001: 5). The antirealists with whom we are concerned are concerned with accommodation – their antirealist proposals are such that they can incorporate received opinion. In many (or even most) areas of discourse, the received opinion is that statements in the domain *are* subject to bivalence, and so unless the antirealist is willing to go against the received opinion, the Dummettian criterion is unsatisfactory.

Fine then claims that for a non-sceptical form of antirealism to be possible (i.e. a form of antirealism not at odds with received opinion), we require a distinction between first-order propositions that state how things are without regards to the metaphysical status of those things, and second-order claims that comment on the status of the first-order propositions, and are independent from them (Fine, 2001: 11). We require, for example, first-order propositions such as <there are tables>, and second order propositions that are about whether <there are tables> is true in the real, objective, mind-independent sense, or just in the everyday sense. These second-order propositions need to be understood as different from the first order ones, in order to block the possibility that the second order proposition might itself not be describing objective reality. For a distinction between realism and antirealism in a given domain of discourse to be maintained, it is necessary for there to be a way to distinguish propositions about what is really the case from the claims of ordinary, non-metaphysical talk.

On Fine's view, sophisticated antirealisms require a metaphysical conception of reality, so that the antirealist can consistently affirm that something is the case whilst denying that it is 'really' the case. Fine (2001: 2) discusses two answers to the question of what this metaphysical conception of reality might consist in. The first identifies metaphysical reality with what is 'objective' or 'factual', such that an antirealist in the relevant domain denies that there are any facts 'out there' making true the propositions in that domain. (For example, the ethical expressivist denies that moral judgements are anything other than expressions of attitudes.) The second way to understand a metaphysical conception of reality is to identify that conception with what is 'fundamental'. The antirealist in a given domain then claims that facts in that domain are reducible to facts of some other sort.

Both of these conceptions of metaphysical reality seem to offer us a distinction between realism and sophisticated antirealism; it is perfectly compatible with affirming some proposition to deny that it is genuinely factual/fundamental (Fine, 2001: 3). But now we can see how the minimalism gets creeping. The antirealist about tables, when she says that this table is bigger than that one, takes herself to be making a claim about how things are, but not how things really are. She thinks it's a fact that this table is bigger than that one, but it's not a FACT that this table is bigger than that one. The block capitals are supposed to represent the metaphysically inflated, mind-independent notions associated with realism,⁵⁸ but an example from Dreier (2004: 28-30) illustrates why this can't quite be right.

Imagine a moral antirealist who agrees with the realist that slavery is wrong, and so agrees that slavery has the property of moral wrongness. The antirealist is unlikely to say that slavery would no longer have the property of wrongness 'if only we slavery-haters mellowed out a little' (29), and so the wrongness of slavery isn't (entirely) mind dependent. It is a property, but not, according to the moral antirealist, a PROPERTY, and so a PROPERTY cannot just be a mind-independent property. (And as Dreier (2004: 30) points out, whilst it is almost tempting to define a PROPERTY as a MIND-INDEPENDENT property, that doesn't look very promising either.) We can give the same treatment to inflated metaphysical notions like REALITY, OBJECTIVITY and FUNDAMENTALITY. In the face of

⁵⁸ This device is borrowed by Dreier from Timmons (1999), who himself adapts it from Putnam.

creeping minimalism, there is no way to articulate a notion of metaphysical reality that floats free from our ordinary notion of reality. The distinction we have been relying on to state the antirealist position breaks down.

These kinds of arguments justify two main quietist positions. The first is that there is no way to state a difference between ordinary and metaphysical discourse, and so the notion of a distinctively metaphysical discourse is devoid of content. There is no higher-order view of how reality is – reality is just what we encounter and represent in thought and language. Fine (2001: 12) takes this type of quietist view to be untenable on the grounds that merely adopting the view presupposes the intelligibility of a metaphysical conception of reality, because it requires taking only the first-order facts as real. (This argument seems a little quick – the quietist need not deny the reality of second order facts in a way that depends on a metaphysical conception of reality. The view can be stated as simply the idea that only ordinary discourse has content, and the quietist can literally remain silent about anything else.) Fine claims further that there is strong intuitive evidence in favour of an intelligible notion of objective reality, and that the need for the notion in framing metaphysical questions makes it almost impossible to deny (2001: 13).

Even if the notion of objective reality is intelligible, Fine identifies a nevertheless troubling form of methodological quietism – there is no way of ascertaining which propositions are about objective reality – ‘judgements...would appear to be metaphysical in the pejorative sense of floating free from any considerations that might tell for or against their truth’ (Fine, 2001: 13). The point is that it is not clear what considerations for or against particular claims about metaphysical reality might be. The notion of objective reality becomes useless because answers to questions about it cannot be found; we can make no progress by enquiring into what metaphysical reality is like.

Fine’s idea for solving this problem is to appeal to grounding. Fine thinks that grounding talk exchanges problematic questions about what exists in metaphysical reality for more tractable questions about what grounds what, and rescues the notions of factuality and fundamentality from the quietest threat. For the sake of simplicity, I’ll focus here on the conception of metaphysical reality as what is objective or factual. What is important for our purposes is that we can say things about what exists in metaphysical reality and distinguish factual from non-factual propositions by appeal to the grounds for

those propositions. I'll explain this with a couple of examples: the first about composition, and the second about morality.

Take the realist and the antirealist about tables, who both agree that this table is bigger than that one. The realist might think that a robust, metaphysical explanation for the fact at hand has to do with the way that particles are arranged such that they constitute tables of different sizes. The antirealist might agree that facts about the arrangement of particles partially ground the table fact, but suppose she thinks that table talk is a myth apt for conveying content that is about the arrangements of particles. Because she thinks table talk is non-representational (in the sense that it does not represent content that is about what we might assume given the surface grammar of the relevant statements), she also thinks that the table fact has some non-factual ground. Amongst the grounds for statements about tables is the table myth. The fictionalist thinks therefore that the proposition <this table is bigger than that one> is non-factual (because it has some non-factual ground), and so is not committed to whatever is quantified over in a Quinean paraphrase of the sentence expressing that proposition. The realist takes the proposition about tables to be factual, and is correspondingly ontologically committed to tables.

Here's the second example. Our ethical antirealist thinks, as does the moral realist, that it's wrong to torture children. But realist and antirealist give differing accounts of what this wrongness consists in, or of what grounds the truth of the proposition (P) that it's wrong to torture children. The realist might say that P is basic (its being the case that P is true consists in nothing more than its being the case that it is wrong to torture children). The antirealist, on the other hand, might be an expressivist, and say that amongst the grounds for P is something to do with speaker-attitudes (it's being the case that P is true consists in nothing more than having some particular non-cognitive attitude towards child torture). Where the realist must make reference to the contested element (in this case wrongness) when citing the grounds for P, the expressivist does not. For the expressivist then, there are at least some constituents of P that are non-representational; there is no wrongness in the world that can make true P. It follows that at least part of the grounds for P are non-factual. Since no factual proposition can have a non-factual ground (Fine, 2001: 19), it follows that the expressivist takes P to be non-factual.

One reservation: Not all expressivists about morality will be willing to grant that it *is* speaker attitudes that make true the proposition that it is wrong to torture children. Some sophisticated expressivists will insist instead that it is the wrongness of child-torture, or perhaps the pain that a tortured child will suffer that makes it the case that it is wrong to torture children. A ‘sophisticated enough’ antirealist will give the same answer to the grounding question as the realist gives, and so appeal to grounding will not help settle matters. I don’t have a clear answer to this sort of concern, but I do think that appeal to grounding treads an attractive line between extreme quietism on the one hand (whereby the distinction between realism and antirealism collapses), and accommodation and minimalism on the other.

There is no obvious reason why antirealists about a given domain of discourse would *want* to answer grounding questions in the same manner as the realist. Sophisticated antirealists are concerned with accommodating ordinary linguistic practice and so it is important that they be able to talk as though sentences in the contested domain are truth apt (and sometimes true), can be believed, and can be embedded in complex linguistic contexts. It is not clear that ordinary linguistic practice extends to answering questions that are explicitly about the grounds for propositions like ‘torturing children is wrong’.⁵⁹ (Though it’s not clear that it doesn’t either. I suspect if a random sample of people were asked what makes it the case that torturing children is wrong, the questioner would be met with blank stares or suspicious frowns, and mumblings about it being ‘just wrong’ or ‘hurting children’ in roughly equal measure.) In any case, there is no clear consensus on what the grounds for the wrongness of child torture are (perhaps it has to do with the law, perhaps with the suffering caused to the child, or perhaps torturing children is wrong because it treats them as mere means rather than as ends in themselves). There is, however, a clear consensus that moral judgements are the kinds of things we can embed in larger linguistic contexts. This makes it far more important that an antirealist have an account of how moral judgements can be embedded than that they have the same account of the grounds of proposition about morality as the realist does.

⁵⁹ This is not in tension with my earlier claims that ordinary speakers use grounding locutions – it is not obvious that ordinary speakers explicitly consider grounds in the theoretical way required for Fine’s proposal to go through.

Suppose, on the other hand, that the sophisticated antirealist does give the very same answers to the grounding question as the realist. Blackburn (2005: 323) claims to recognise that the ‘quasi’ in quasi-realism might be misleading; it suggests that quasi-realists talk ‘as if’ there are elements of the contested domain (e.g. moral truths) even though there are none really. He says:

‘if the words retain an uncorrupted, English, sense then...the quasi-realist holds not just that we talk and think as if there are [moral truths]...but that there are.’
(Blackburn, 2005: 323.)

‘I do not say that we can talk as if kicking dogs is wrong when ‘really’ it isn’t wrong. I say that it is wrong (so it is true that it is wrong, so it is really true that it is wrong, so this is an example of a moral truth, so there are moral truths).’ (Blackburn, 2001: 319.)

Blackburn’s comments suggest a rejection of the way we (following Fine) have set things up – to speak of a metaphysical reality that the antirealist thinks sentences of the contested domain fail to represent is to make a kind of mistake. To evaluate this kind of position would be to go beyond the scope of the thesis, but I think a defender of this sort of view is unlikely to find creeping minimalism too much of a problem. They won’t share the intuition that minimalism sucks the substance out of weighty metaphysical notions, or that the propositions of the antirealist are imposters. Appeal to ground cannot distinguish them from the realist, because their position does not correspond to the distinctions involved in the setup of the problem grounding is being invoked to solve. For this reason, I’ll set this position aside.

If we like, we can now state a new ‘Finean’ criterion of ontological commitment. One is ontologically committed to whatever is the value of a bound variable in a translation into their canonical form of the sentences expressing the propositions one takes to be factual. The expressivist who takes propositions about wrongness to have (at least partially) non-factual grounds and thus to be non-factual is not ontologically committed to a property of wrongness, whereas the realist who thinks such propositions have entirely factual grounds, is.

If successful, Fine's proposal for distinguishing realism from sophisticated antirealism is itself an important theoretical use for grounding talk, and therefore a motivation for continuing to talk in terms of grounding. In other words, it is a reason to prefer antirealism over eliminativism about grounding talk (once we are convinced we shouldn't be realists). But one might worry that Fine's proposal won't work if we are antirealist about grounding itself. This is the topic of section 7.3.

7.3. Sophisticated Antirealism about grounding

In sections 4 and 5 I developed models for antirealism about grounding. Appeal to ground to distinguish realism from sophisticated antirealism can only be used to motivate continued use of grounding locutions by an antirealist about grounding if they can sustain the relevant distinction despite their antirealism. I'll begin by working through an example to demonstrate that the antirealist about grounding can legitimately appeal to grounding to make the distinction. I'll then address the question of how to distinguish a realist about grounding from a sophisticated antirealist about grounding.

Suppose Amy is a (sophisticated) moral antirealist, and Richard is a moral realist. Both Amy and Richard say P: 'torturing children is wrong'. In order to ascertain their different ontological commitments, we ask Amy and Richard about the grounds for their positions. Suppose Richard thinks that P is basic; it is just the wrongness of child-torture that makes true the proposition that it is wrong to torture children. Richard thus makes reference to moral wrongness when he gives his account of the grounds for P. Amy does not, and so we can now tell that since Richard, but not Amy, thinks that P is a factual proposition, Richard, but not Amy, is committed to the existence 'in reality' of moral properties.

For any domain of discourse about which one can be an antirealist, by the processes of accommodation and minimalism one can become a *sophisticated* antirealist. An antirealist about grounding hoping to accommodate the way in which grounding claims are made (both by the folk and in philosophical theorising) will hope not only to assert grounding claims, but also to secure the right to say that those claims are true, to embed them in more complicated linguistic contexts, and so on. Once we have the possibility of sophisticated antirealism about grounding, minimalism gets creeping again. How are we then to distinguish realist from sophisticated antirealist *about grounding*?

Before we answer that question, it is worth getting clear on just what an antirealist about grounding takes to be going on when realist and sophisticated antirealist in some domain of discourse disagree about the grounds for some proposition. I'll use expressivism on the credence model as an example, but I hope it will be clear how alternative antirealist accounts will deal with the relevant claims. When Amy (an expressivist about morality) says that P is grounded in her con-attitude towards child torture, the expressivist about grounding takes her to be expressing her greater willingness, in the light of new evidence, to update or revise her belief that she has a con-attitude towards child torture than to update or revise her belief that torturing children is wrong (where those beliefs are connected in an explanatory way in her web of belief). When Richard (a realist about morality) says that P is basic, the antirealist about grounding takes him to express his unwillingness to update or revise his belief that torturing children is wrong in the light of new evidence. Richard is inclined to treat his belief as highly resilient. (Of course, the realist need not say that P is basic; perhaps he thinks P is grounded in the suffering that torturing children causes those children. In that case, the antirealist takes him to express his greater willingness, in the light of new evidence, to update or revise his beliefs about the suffering child-torture causes children than to update or revise his belief that child torture is wrong.)

The antirealist about grounding takes the realist and the antirealist about morality to express different attitudes with their different grounding claims. This is enough to distinguish realist from sophisticated antirealist about morality, and so to do the work that grounding was supposed to do here. As per Fine's original proposal, we can determine whether a proposition is factual or non-factual according to the different parties to the debate on the basis of the story they tell about how that proposition is grounded.

Satisfied that an antirealist about grounding can still appeal to grounding to distinguish between realist and sophisticated antirealist, we can now turn to the question of how to distinguish realism about grounding from sophisticated antirealism about grounding. Suppose Steph is a sophisticated antirealist both about morality and about grounding. Steph believes that moral claims like P (that torturing children is wrong) are grounded in attitudes. Let's suppose she is a Yablo-style fictionalist about grounding, holding that grounding talk is a useful game of make-believe apt for conveying information about comparative degrees of confirmation enjoyed by different propositions in a theory, and the entailment

relations between those propositions. Steph's friend Gordon is also an expressivist about morality, but he is a realist about grounding. That is, he believes that there is an objective, mind-independent notion of ontological dependence holding between moral claims and attitudes, such that the latter provide a robust, metaphysical explanation for the former. Steph and Gordon assert the very same sentences, despite their seemingly differing worldviews.

Both Steph and Gordon assert Q: 'moral claims are grounded in attitudes'. Suppose we now ask Gordon 'what are the grounds for your claim Q?' Gordon thinks Q is a factual proposition, because it has no non-representational components (of course, Gordon's realism about grounds for moral claims doesn't commit him to realism about moral *properties*). We should expect Steph to give grounds for Q that include some non-representational components (as is constitutive of her antirealism about grounding). Gordon, but not Steph, should think that the proposition <moral claims are grounded in attitudes> is a factual proposition. But the issue is complicated by the fact that we are trying to characterise Steph's antirealism by appeal to the very thing about which she is antirealist. Perhaps if Steph's antirealism about grounding were sophisticated enough she might not only assert Q, but also give the same response as Gordon gives to the question of how Q is grounded.

Note that this case is subtly different from the discussion at the end of section 7.2. There, the issue was that a quasi-realist like Blackburn might reject the characterisation of the debate between realists and antirealists as being about representing metaphysical reality. Steph could be happy with a conception of metaphysical reality, deny that grounding claims are in the business of stating facts about it, but nevertheless give the same response to the grounding question as Gordon. Her sophisticated antirealism about grounding licenses her to assert the same sentences *about grounding* as the realist asserts.

Let's suppose this is the case. When we ask Steph about the grounds for Q, her fictionalist interpretation of the question is something like 'which proposition is it that you take to be comparatively better confirmed than, and to entail the proposition that (propositions about attitudes are comparatively better confirmed than the propositions about moral claims that they entail)?' Steph might answer that it is just the bracketed proposition in her interpretation of the question. In other words, it is just 'moral claims are grounded in attitudes' (the conveyed content of which is that propositions about attitudes are

comparatively better confirmed than the propositions about moral claims that they entail. But of course, Steph's answer to the question 'what grounds Q' ('moral claims are grounded in attitudes') is precisely the answer that will be given by Gordon, who takes the proposition <moral claims are grounded in attitudes> to be factual; it has no non-representational constituents.

Were this Steph's answer, we would be entirely unable to distinguish realist from sophisticated antirealist about grounding. Steph and Gordon would give all the same answers to iterations of questions about grounding grounding claims. Granting the possibility of sophisticated antirealism about grounding, this would constitute (at the very least) a limitation of Fine's proposal such that it is not the case that questions of realism in any domain can be settled by appeal to ground. But I'm inclined to think this shouldn't be Steph's answer.

Steph has earned herself the right to engage in grounding talk as the realist does. She thinks grounding talk is useful and worthwhile, and she has a story about what is going on when she makes a grounding claim. But what she thinks about grounding claims is different from what the realist thinks. Steph thinks that what grounds grounding claims does not have to do with grounding, but rather has to do with attitudes and inferential connections between beliefs, or with conformational structures. Since Steph's sophisticated antirealism allows her to use grounding locutions to express her position, she should be able to make clear her attitude to grounding.

Recall that questions of ground can be answered whilst taking a metaphysically neutral stand on whether or not the grounded entities exist. When the sophisticated antirealist about ethics is asked 'what grounds the proposition <torturing babies is wrong>?' she does not interpret the question to be 'what grounds the fact that you have a con-attitude towards torturing babies?' Rather, the grounding question asks her to say what it is she takes to account for the truth of the proposition <torturing babies is wrong>. Similarly, when Steph is asked what the grounds are for her assertion Q, she is being asked to use grounding locutions to talk about grounding. She can do so whilst taking a metaphysically neutral stand on the existence of the grounded entities (in this case, the constituents of the proposition Q). Steph thinks that what makes it true that moral claims are grounded in attitudes is (for example) that

moral claims are less well-confirmed than the attitudes that entail them. This then is her answer to the question ‘what grounds Q?’

Once we have distinguished Steph and Gordon’s views by taking into account the different grounds they cite for Q, there is no further proposition to which they will both assent, and about which we do not have a way to distinguish views that seem to differ. Steph will not assent to Gordon’s claim that it is a factual notion of grounding that grounds Q, and Gordon will deny Steph’s contention that the grounds of Q have to do with attitudes or credences. If Steph takes Q to be non-factual, she will also think questions about what grounds the non-factuality of Q will have non-factual answers; similarly for Gordon. For any questions about grounding to which the same response is given by realist and sophisticated antirealist, we can ascend one level up and ask another iteration of the grounding question (i.e. ‘what grounds S?’ where S is the proposition realist and sophisticated antirealist both assert). Any resultant regress is benign, because a distinction between realist and sophisticated antirealist can be made wherever it is required, and higher order realism or antirealism about grounding does not affect use of grounding idioms in making lower-order distinctions.

8. Conclusion

Antirealism about grounding is motivated by a need to respond to the grounding-explanation problem discussed in the previous chapter, and by concerns about the epistemology of both specific grounding claims and claims about grounding. Unlike eliminativism about grounding talk, antirealism about grounding licenses us to continue to appeal to grounding to do the useful work in metaphysics for which many have argued it is well suited. I outlined versions of both expressivism and fictionalism about ground, each of which could function as the basis for a more comprehensive theory. As demonstrated in section 7, grounding on an antirealist conception is capable of carrying out the role that many have taken grounding to play in distinguishing realism from sophisticated antirealism. In my concluding remarks, I’ll make explicit some connections between antirealism about grounding, explanation, and reality’s structure.

Concluding Remarks: Antirealism, Nonsymmetry, Explanation and Structure

Of the seven strategies I identified for characterising the notion of fundamentality at the end of Chapter 1, the most widely discussed are those which appeal to naturalness or to Sider's 'structure', and those which appeal to grounding. In the second chapter I argued that the property *being perfectly natural* falls significantly short of playing the role of a perfectly natural property, and so that characterising the structure of reality by appeal to naturalness would undermine the metaphysical significance of that structure. In a similar vein, I argued that Sider's primitive notion of absolute structure cannot do the important work in metaphysics that Sider takes to justify positing it. The way in which structure underwrites the explanatory utility of structural notions cannot account for the explanatory role of absolute structure. Neither Sider's structure nor Lewisian naturalness are suited, by their own lights, to characterising the structure of reality.

Having rejected two of the most promising three strategies for characterising the structure of reality, in the second part of the thesis I took up discussion of the notion of grounding. I think there are some problems with the orthodox conception of grounding (insofar as there is any orthodox conception) and I explored them in chapters 3 and 4. In chapter 3 I challenged the common assumption that grounding is an asymmetric relation. An asymmetric conception of grounding has trouble accommodating some intuitive cases of dependence that can be adequately described if grounding is instead nonsymmetric. The view I there developed, metaphysical interdependence, is, I argued, a position worth taking seriously in the debate over metaphysical structure. Not only does it accommodate otherwise irreconcilable intuitions about grounding, but it is more theoretically virtuous and consistent with more of the ways reality might turn out than rival theories are.

My second challenge to the orthodox conception of grounding concerned the purported link between grounding and metaphysical explanation to which proponents of grounding often appeal. It is very common to find friends of grounding hoping to elucidate the somewhat problematic notion of grounding by relating it to the more familiar notion of explanation. I argued that in order for this to be viable, the relevant notion of explanation must be one of successful explanation, which is influenced by agent-relative features. The ideas that grounding and explanation share a connection as close as has been claimed in the literature, and that the relevant notion of explanation is one that is affected by agent-relative considerations are jointly in conflict with the idea that grounding is a feature of objective, mind-independent reality.

My solution to this apparent tension was to reject the thesis that grounding is a feature of mind-independent reality, and instead to motivate and develop antirealist approaches to grounding claims. So long as we appeal to grounding to limn the structure of reality, antirealist conceptions of grounding make that structure a mind-dependent feature, rather than a feature of the mind-independent world. Grounding relations are explanatory relations because they communicate information about features of our cognitive (epistemic or conceptual) architecture. What makes for a successful explanation is affected by what stands in need of explanation, and what counts as explanatory for an agent.

Antirealist approaches to grounding are able to accommodate not only the explanatory nature of grounding claims, but also the nonsymmetry for which I argued in Chapter 3. Some pair of beliefs might occupy symmetrical and equally central roles in web of beliefs, or some pair of propositions prove equally intelligible and mutually supporting.⁶⁰ Such cases are represented with grounding claims that are structured in accordance with metaphysical interdependence.

Antirealism about grounding offers us a means to reconcile our concern with the structure of the world with our failure to legitimately characterise the metaphysical structure of the world through extant realist approaches. We care about those aspects of reality about which we have beliefs in which we have robust confidence, or which are central to our theoretical or conceptual schemes. In other words, we care about

⁶⁰ As grounding is a one-many relation, the grounding side of this pair might itself be a set of beliefs or of propositions.

what helps us to most effectively understand the world, and it is these features which we project onto reality when we describe its structure using grounding talk.

Appendix: Reference Magnetism

The Lewisian doctrine of reference magnetism is an account of reference which appeals to the natural properties (see Chapter 2). This aspect of the naturalness role is somewhat unusual in that an application of fundamentality to reference is rarely discussed outside of the context of naturalness (or Sider's structure). In conversation, I have found it to be fairly common amongst philosophers (albeit not often recorded in print) to take reference magnetism to be the best reason to believe in naturalness theory, and so this appendix is included so as to include a more detailed discussion of reference magnetism than was appropriate in the main text. Reference magnetism is an account of how our representational expressions get their contents, and so is instrumental in solving a number of difficult problems. In this appendix I describe the main problem that Lewis develops reference magnetism in an attempt to solve, and discuss that proposed solution. I also present some challenges to the doctrine of reference magnetism.

1. Putnam's paradox

Instrumental in the development of Lewis' naturalness theory is his response to a question posed by Putnam concerning reference. 'Putnam's paradox' is Lewis' name for a problem discussed by Putnam (1977), but that also famously plagued Quine (e.g. 1960; 1969) and Davidson (e.g. 1984). The problem is that there is no 'semantic glue' to stick our words (or our thoughts) to their referents apart from our intentions to refer in such a way that what we say comes out as true. I'll begin by describing the paradox.

Metasemantic theories are theories about the nature of semantic facts; they tell us in virtue of what a name refers to something, or in virtue of what a predicate applies to some object). Lewis's concern is with 'interpretationist' metasemantic theories. As characterised by Williams, 2007: 363, such theories proceed in two steps: (1) they match up sentences with a range of appropriate candidate interpretations of those sentences, and (2) they fix the correct interpretation in accordance with the best theory of the world. One such interpretationist theory, and the one I will focus on here, is 'global descriptivism', which

is described below (though the criticisms and responses developed should be taken to apply more widely).

The global descriptivist's first step is to construct a theory. For any given domain, this is a matter of gathering together platitudinous sentences in that domain. In the domain of electron science, for example, the theory would include expressions such as 'electrons have a negative charge', 'electrons form clouds around the nuclei of atoms', and 'electrons have a mass of approximately 9.1×10^{-31} kg'. The theoretical term 'electron' is then defined implicitly *by* the theory through the process of ramsification (i.e. through replacing all the theoretical terms of the theory with higher order variables, bound by higher-order quantifiers (see Lewis, 1970)).

Total theory is the sum of all of the platitudes in every domain, and contains all of the sentences that ordinary speakers would assent to. Global descriptivists identify a semantic theory that renders as many as possible of the sentences of the total theory true. If global descriptivism is true, then every expression gets its referent based wholly on the descriptions associated with that expression – the referent of 'electron' is whatever satisfies the descriptions associated with 'electron', the referent of 'atom' is whatever satisfies the descriptions associated with 'atom', and so on.

According to global descriptivists, the only constraints on the referent of an expression are internal to the descriptivist theory, since they are based entirely in the dispositions and beliefs of the user(s) of the term (Chalmers, manuscript: 2). It is this that makes room for Putnam's objection to global descriptivism. Putnam's claim is (in essence) that any world can satisfy any theory (Lewis, 1984). For any way that the world is, there will be a reference scheme (an assignment of referents to expressions) that makes any (total) theory come out as true. Putnam's objection demonstrates that the patterns in linguistic usage can be accounted for with radically deviant assignments of referents to expressions, rendering the reference relation *indeterminate*.

Consider an adaptation of an example from Sider (2011: 24). We have a set of sentences S such that our words refer to whatever they must in order for S to come out as true (this is global descriptivism). Sider asks us to consider the intuitively false sentence 'some pigs have wings'. This sentence will be false according to some interpretation iff nothing that falls under the extension 'pig' also falls under the

extension 'has wings'. If, according to some interpretation, there is something that falls under the extension of both terms, then our sentence will be true according to that interpretation. The only thing constraining reference is our intention to refer in such a way that the sentences we take to be true come out as true, and the sentences we take to be false come out as false.

So, it is possible to assign crazy interpretations to the terms of our language. Suppose that instead of assigning the set of pigs to 'pigs', we instead assign 'pigs' the set of hard boiled eggs. To 'has wings' we assign the set of things that have runny yellow insides. The sentence 'some pigs have wings' correctly comes out as false, even though the referents of the expressions involved are neither the set of pigs nor the set of things with wings. If S also contains the sentence 'every pig is an animal', our crazy interpretation must make this sentence come out as true. It can do so by, for example, assigning 'is an animal' to 'is an egg'. Since everything that is a hard-boiled egg is an egg, 'all pigs are animals' will indeed come out as true.

Consider a further example. Assign the referent of 'London' to Sydney, and the referent of 'Sydney' to London. In order to make sentences such as 'London is in England' and 'Sydney is in Australia' come out as true, we must assign 'England' to Australia and 'Australia' to England. When we now say 'it is warm in Australia in January', in order to compensate for the assignment of Australia to England we must take the set of cold things to come under the predicate 'is warm' so that the above sentence comes out as true under our interpretation. Crazy assignments of referents to names are cancelled out by crazy assignments of extensions to predicates, so that the truth values of the sentences of S are unaffected.

This method can be extended to apply to an entire natural language. In this way we can construct multiple interpretations that make the sentences of S come out as true, and we have no way to distinguish between correct and crazy interpretations. More than this, there is *no fact of the matter* whether the interpretation assigning 'London' to London or the one assigning 'London' to Sydney is the correct one – reference is indeterminate. David Lewis (1984) seeks a way out of this troubling conclusion. His strategy is to argue against the claim that there are multiple competing interpretations of a language between which reference is indeterminate. There must, says Lewis, be some additional constraint on reference that is external to the beliefs and attitudes of the speaker – external to the total theory.

Note that this is a possibility that Putnam himself considers and rejects. According to Putnam, any further constraints will be ‘just more theory’. The thought is that any candidate additional constraint on reference, if it seems good, will be incorporated into total theory. Take, for example, a causal constraint on reference (such that reference is constrained by ‘causal chains that lead into the referrer’s head from the external things that he refers to’ (Lewis, 1984: 226). Any intended interpretation of total theory must then make the causal constraint i.e. ‘reference is constrained by chains that lead into the referrer’s head from the external things that he refers to’, along with the rest of total theory, come out as true. It will still be true that any world can satisfy any theory, including a theory that includes the causal constraint. The casual constraint is ‘just more theory’ – it is subsumed into total theory and the problem of inscrutability applies in the same way as it did previously.

Lewis rejects Putnam’s claim that any further constraint on reference will fall into the above trap (1984: 255). He argues that the additional constraint is not to be imposed by simply accepting the additional constraint and adding it to total theory. Instead, the additional constraint requires the interpretation of the total theory to *conform* to that constraint. The constraint is not *part* of the theory, but is a constraint *on* the interpretation of the theory.

Lewis’ answer to the problem of radical inscrutability (to Putnam’s Paradox) is to take the referents themselves to play a role in determining the reference of our words. He claims that most of the things that there are are ‘miscellaneous, gerrymandered, ill-demarcated’, and that ‘only an elite minority are carved at the joints, so that their boundaries are established by objective sameness and difference in nature’ (1984: 227). It is this elite minority, the perfectly natural properties, that are most eligible to serve as the referents of our expressions. What might previously have counted as other potential candidates are rejected before they are even considered on the grounds that they are too unnatural.

One worry here is that it is not entirely clear how it could be the case that the additional constraint is not part of the theory. Total theory is, as we saw above, the collection of all of the sentences assented to by ordinary speakers (and, presumably, experts in a given technical field). Theories about semantics are themselves part of total theory, and so it seems right that whatever we say about interpreting total theory

must also apply to the metasemantic proposal under consideration. Reference magnetism is a metasemantic proposal.

Perhaps this objection can be avoided by distinguishing between the first order language and a metalanguage, in order to use the metalanguage to make metasemantic claims that constrain the first order total theory. But this seems quickly to lead to a problem in interpreting the metalanguage. In order to constrain reference in the metalanguage we will need a third order language, and so on ad infinitum. This is a problematic regress because it means that no constraint would ever actually be realised – in order to constrain the first order language we must be able to make claims in the metalanguage, to be able to make claims in the metalanguage it must be constrained by a metametalanguage, and so on.

So maybe the Lewisian thought is that the metasemantic theory is not itself to be evaluated using semantics. It is simply a fact that reference is constrained by something external to the intentions of the speakers, and this particular fact is true independently of speakers' intentions or anything else. The sentence 'eligibility constrains reference' is part of total theory and gets treated as such, but the constraint it expresses is supposed to be a constraint *on* the theory. There is a lot of interesting literature on Putnam's 'just more theory' manoeuvre and its application to constraints on reference (see e.g. Button, 2013) but it is beyond the scope of this thesis. For the sake of argument, I will assume that Lewis' constraint does not fall foul of the 'just more theory' objection, and evaluate it on different grounds.

2. Reference magnets and fundamental physics

On Lewis' proposal, the correct interpretation of our language must, as far as possible, assign natural properties and relations or their extensions to the predicates of the language. Note again that it is vital that the predicates actually stand (as far as possible) for the natural properties and relations, not merely that something like 'predicates stand as far as possible for natural properties and relations' come out as true on the interpretation – that would be to fall into the 'just more theory' trap. The more natural the property, relation, or assignment of an extension to a predicate, the more eligible that assignment. As Lewis notes, since eligibility is therefore a matter of degree, making the total theory come out as true will also be a matter of degree (Lewis, 1984: 228).

The most eligible referents are, Lewis stipulates, the fundamental physical properties (1984: 228). One concern with the idea that the fundamental physical properties are deserving of elite reference magnetic status is that we have only very recently had words for such properties, but have had words for everyday macroscopic objects for far longer. Lewis raises this problem, and presents us with a solution. His answer is that we used to lack words for joint carving notions because the correct interpretations of our language were ‘the ones that did best on balance, not the ones that did best at best’ (Lewis, 1984: 228). Physics discovers the most elite classes, but other classes are also fairly elite because they can be defined fairly easily in terms of the fundamental physical notions.

Taking for granted that the perfectly natural properties are those discovered by fundamental physics, it is still far from clear that this fact makes those properties the most eligible referents of our language. Though the success of science might seem to suggest that the terms of our language attach most readily to the perfectly natural properties (i.e. we *are* able to successfully represent the world at the level of fundamental physics, and so we are warranted in looking for a metasemantic theory that explains this success) I don’t think the success of science provides an argument that the properties of fundamental physics are the most natural referents of our language.

As science progresses, so does language, and the meanings of our terms shift. It is not the case that whenever (or, at least, often) when we intend to use an expression to refer to fundamental physical properties we are successful. In fact, it would be detrimental to science if we were. Suppose we identify a role for some fundamental physical property and label the property ‘ λ ’. There is in fact nothing that plays the λ role as we understand it, but there *is* some extremely natural property in the vicinity. This property is magnetic enough to attract the reference of our expression ‘ λ ’, even though our use of ‘ λ ’ is such that not all of the things we say about λ can come out as true, since ‘ λ ’ in fact picks out the extremely natural property in the vicinity of the role we thought λ played. The upshot is that we end up referring to a property that plays a *different* role to the role we believe it plays, and so end up believing false things about that property. If reference magnetism allows us to get away with having false beliefs about the fundamental properties, it seems that the success of science does not provide an argument for reference magnetism. If reference magnetism is true, then science is successful despite the fact that our scientific

theorising concerns properties that play different roles than the roles we take them to play in our theories.

3. Use and Eligibility

The reference magnetist holds that the way in which expressions are used in a language constrains the possible interpretations of that language (recall that, as far as possible, an interpretation of a language ought to make the sentences of the global folk theory come out as true on that interpretation). The eligibility of the referent then decides between competing interpretations. Reference magnetism is therefore characterised by a trade-off between the eligibility of the referent (which is, on Lewis' view, entirely independent of the speakers of the language) and charity to use. In line with Lewis, Sider (2011: 31) defines charity to use as 'the number (on some suitable measure) of sentences believed...by the speakers of the linguistic community that come out as true under that interpretation'. The trade-off between use and eligibility generates some indeterminacy, but it will be far less pervasive than is the case when there is no externalist constraint.

It is a strange consequence of reference magnetism that speakers of a language might be in the dark about the referents of their own expressions, as they do not have access to the facts about which, out of the competing interpretations that are consistent with charity to use, is the most eligible. It seems important not just that our words *do* refer determinately, but that we have some mechanism for discovering what the determinate referents of our expressions are.

The problem is that any access we have (or take ourselves to have) to the facts about eligibility generate theory *internal* constraints on reference. Truly external constraints are constraints that are opaque to us. If naturalness theory is to be adopted, then it should form part of our core theory in a given domain. Platitudes about e.g. *mass* would include '*mass* is a perfectly natural property'. Reference magnetism dictates that we find the interpretations of the language such that the platitudes of global folk theory -- which includes platitudes about the naturalness of *mass* -- come out as true on those interpretations. Of course, if reference magnetism is true then it is *also* the case that the eligibility of *mass* is making it easier for us to refer to *mass* than to other properties in the vicinity. But the point is that the work the eligibility

of *mass* is doing to fix the referent of our expression ‘mass’ is entirely independent from our theorising. In that case, we are in no position to appeal to it in order to discover how our expressions do in fact refer.

Consider the following example. Ordinary English speakers often assent to the following sorts of sentences:

- (1) The only fruit I like is apple, but my favourite food is tomato
- (2) I’ve eaten no fruit today, but I had three types of vegetable for dinner and a cheese and tomato sandwich for lunch
- (3) I’m not growing any fruit this year, just cucumbers and tomatoes

A speaker of English unsure whether or not the predicate ‘is a fruit’ properly includes tomatoes in its extension might take the readiness with which speakers assent to (1)-(3) to establish that the extension of the predicate ‘is a fruit’ does not include tomatoes. But suppose our English speaker is impressed by reference magnetism, and so she takes her discovery about the use of the predicate to be just one side of the story. Her next step is to try to establish the eligibility of the referent that her study of use had turned up for the expression ‘is a fruit’. She further examines platitudes surrounding fruit, and determines that the biological kind **fruit** is more natural than the kind **fruit minus tomatoes**.

Our speaker might take her discovery to indicate that the property *being a fruit* (call it *F*) is more natural than the property *being a non-tomato fruit* (call it *T*), and therefore to think that the referent of the predicate ‘is a fruit’ is property *F*. But nothing about this is theory external (the results of her investigation might have gone the other way, so that she took ‘is a fruit’ to pick out property *T*). Our speaker will now adjust her use of the expression ‘fruit’ in accordance with her apparent discovery, and perhaps eventually platitudes surrounding fruit will change. It is our speaker herself who judges the degree to which the properties in question are natural.

There are two options here. We might think that this is the last word, and so reference magnetism fails to provide any external constraint on reference. Speakers of a language might adjust the use of their expressions in accordance with their judgements about naturalness, and this might in turn affect the meaning of those expressions, but only because it impacts on the use of the expressions. Chalmers

(manuscript: 3) calls this *weak* reference magnetism. Contrast this with *strong* reference magnetism, where the eligibility of the referent really does play a theory-external role (independent of what speakers of the language believe about the eligibility of particular interpretations).

Speakers who are mistaken about naturalness judgements will, if naturalness really does play a theory-external role, make systematic and undetected errors about what their terms mean. This becomes more troubling when we consider questions about, for example, whether definitions such as ‘mares are female horses’ can be trumped by naturalness considerations. If they can, we might have very little idea about what the terms of our language refer to and thus an extremely unexplanatory and counterintuitive theory of reference. If naturalness cannot play a trumping role in such cases then reference magnetism is not a global metasemantic thesis, but rather operates like local descriptivism so that the reference of some expressions is fixed by description and the reference of others by naturalness (Chalmers, manuscript: 5.)

My view is that reference magnetism is a novel approach to reference, but one that does not quite have the resources required to overcome the problems with constraining reference (only one of which I have discussed here). If my arguments here are successful, they further undermine the motivation for naturalness theory, and the hope of appealing to naturalness to characterise the structure of reality.

References

- Aristotle (1998) **The Metaphysics**. Translated by Lawson-Tancred, H. London: Penguin Books Ltd.
- Armstrong, D. M. (1997) **A World of States of Affairs**. Cambridge: Cambridge University Press.
- Armstrong, D. M. (2004) **Truth and Truthmakers**. Cambridge: Cambridge University Press.
- Audi, P. (2012a) "A Clarification and Defense of the Notion of Grounding." In Correia, F. and Schneider, B. (eds.) **Metaphysical Grounding: Understanding the Structure of Reality**. Cambridge: Cambridge University press. pp. 101-121.
- Audi, P. (2012b) Towards a Theory of the In-Virtue-Of relation. **The Journal of Philosophy**, 109: 685-711.
- Baker, A. (2013) Simplicity. In Zalta, E. (ed.) **The Stanford Encyclopedia of Philosophy** (Fall 2013 Edition). [Online] Available from:
<http://plato.stanford.edu/archives/fall2013/entries/simplicity/> [Last accessed July 2014].
- Barker, S. (2007) **Global Expressivism: Language Agency Without Semantics, Reality Without Metaphysics**. University of Nottingham ePrints [Online] Available from:
<http://eprints.nottingham.ac.uk/696/> [Last accessed July 2014].
- Barnes, E. (2012) Emergence and Fundamentality. **Mind**, 121 (484): 873-901.
- Bennett, K. (2011) By our Bootstraps. **Philosophical Perspectives**, 25 (1): 27-41.
- Bennett, K. (forthcoming) (2014) **Making Things Up**. Oxford: Oxford University Press.
- Blackburn S. (1993) Morals and Modals. **Essays in Quasi-Realism**. Oxford: Oxford University Press.
- Blackburn, S. (1984) **Spreading the Word**. Oxford: Oxford University Press.
- Blackburn, S. (2001) **Ruling Passions**. Oxford: Oxford University Press.
- Blackburn, S. (2005) "Quasi-Realism No Fictionalism." In Kalderon, M. (ed.) **Fictionalism in Metaphysics**, Oxford: Clarendon Press. pp. 322-337.
- Bliss, R. (2011) **Against Metaphysical Foundationalism**. PhD Thesis, University of Melbourne.
- Bohn, E. (2009a) An Argument Against the Necessity of Unrestricted Composition. **Analysis** 69: 27-31.
- Bohn, E. (2009b) Must There be a Top Level? **The Philosophical Quarterly**, 59 (235): 193-201.
- Bovens, L. and Hartmann, S. (2003) **Bayesian Epistemology**. Oxford: Oxford University Press.
- Brandom, R. (2000) **Articulating Reasons: An Introduction to Inferentialism**. Cambridge: Cambridge University Press.
- Burge, T. (1986) 'Individualism and Psychology' **Philosophical Review**, 95: 3-45.
- Cameron, R. (2008) Turtles All the Way Down: Regress, Priority and Fundamentality. **The Philosophical Quarterly**. 58 (230): 1-14.

- Cameron, R. (2010) From Humean Truthmaker Theory to Priority Monism. **Noûs**, 44 (1): 178-198.
- Chalmers, D. (manuscript) Twentieth Excursus: Reference Magnets and the Grounds of Intentionality. [Online] Available from: <http://consc.net/books/ctw/excursus20.pdf> [Last accessed July, 2014].
- Chrisman, M. (2012) Epistemic Expressivism. **Philosophy Compass**, 7 (2): 118-126.
- Clark, M. and Liggins, D. (2012) Recent Work on Grounding. **Analysis**, 72 (4): 812-823.
- Corkum, P. (2008) Aristotle on Ontological Dependence. **Phronesis** 53 (1): 65-92.
- Correia, F. (2004) Husserl on Foundation. **Dialectica**, 58 (3): 349-367.
- Correia, F. (2008) Ontological Dependence. **Philosophy Compass**, 3 (5): 1013-1032.
- Correia, F. and Schnieder, B. (2012) **Metaphysical Grounding: Understanding the Structure of Reality**. Cambridge: Cambridge University Press. pp. 1-36.
- Daly, C. Scepticism about Grounding. In Correia, F. and Schneider, B. (eds.) **Metaphysical Grounding: Understanding the Structure of Reality**. Cambridge: Cambridge University Press. pp. 81-100.
- Davidson, D. (1984) **Inquiries into Truth and Interpretation**. Oxford: Clarendon Press.
- deRosset, L. (2010) Getting Priority Straight. **Philosophical Studies**, 149 (1): 73-97.
- deRosset, L. (2013) Grounding Explanations. **Philosophers' Imprint**, 13 (7): 1-26.
- Demopoulos, W. and Friedman, M., (1985) Bertrand Russell's The Analysis of Matter: Its Historical Context and Contemporary Interest. **Philosophy of Science**, 52: 621-639.
- Dorr, C. and Hawthorne, J. (2013) "Naturalness." In Bennett, K. and Zimmerman, D. (eds.) **Oxford Studies in Metaphysics Volume 8**. Oxford: Oxford University Press. pp. 3-74.
- Dreier, J. (2004) Metaethics and the Problem of Creeping Minimalism. **Philosophical Perspectives**, 18: 23-44.
- Dreier, J. (2009) Relativism (and Expressivism) and the Problem of Disagreement. **Philosophical Perspectives**, 23 (1): 79-110.
- Duhem, P. (1962) *The Aim and Structure of Physical Theory* (New York: Holiday House).
- Dummett, M. (1978) **Truth and Other Enigmas**. London: Duckworth.
- Eddon, M. (2013) "Fundamental Properties of Fundamental Properties." In Bennett, K. and Zimmerman, D. (eds.) **Oxford Studies In Metaphysics Volume 8**. Oxford: Oxford University Press. pp. 75-103.
- Eklund, M. (2005) Fiction, Indifference, and Ontology. **Philosophy and Phenomenological Research**, 71: 557-79.

- Eklund, M. (2011) 'Fictionalism' In Zalta, E. (ed.) **The Stanford Encyclopedia of Philosophy** (Fall 2011 Edition). [Online] Available from: <http://plato.stanford.edu/archives/fall2011/entries/fictionalism> [Last accessed July 2014].
- Esfeld, M. (1998) Holism and Analytic Philosophy. **Mind**, 107: 365-380.
- Esfeld, M. (2001) **Holism in Philosophy of Mind and Philosophy of Physics**. Dordrecht: Kluwer.
- Esfeld, M. and Lam, V. (2006) Moderate Structural Realism about Space-Time. **Synthese**, 160: 27-46.
- Evans, Matthew (2012) "Lessons from Euthyphro 10a–11b." In Inwood, B. (ed.) **Oxford Studies in Ancient Philosophy Volume 42**. New York: Oxford University Press. pp. 1-38.
- Field, H. (1980) **Science Without Numbers**. Princeton: Princeton University Press.
- Fine, K. (1994) Essence and Modality. **Philosophical Perspectives**, 8: 1-16.
- Fine, K. (1995a) Ontological Dependence. **Proceedings of the Aristotelian Society**, New Series, 95: 269-290.
- Fine, K. (1995b) "Senses of Essence." In Sinnott-Armstrong, W., Raffman, D., and Asher, N. **Modality, Morality and Belief: Essays in Honour of Ruth Barcan Marcus**. Cambridge: Cambridge University Press. pp. 53-73.
- Fine, K. (2001) A Question of Realism. **Philosopher's Imprint**, 1 (1): 1-30.
- Fine, K. (2010) Some Puzzles of Ground. **Notre Dame Journal of Formal Logic**, 51 (1): 97-118.
- Fine, K. (2012a) The Pure Logic of Ground. **Review of Symbolic Logic**, 5 (1): 1-25.
- Fine, K. (2012b) "A Guide to Ground." In Correia, F. and Schneider, B. (eds.) **Metaphysical Grounding: Understanding the Structure of Reality**. Cambridge: Cambridge University Press. pp. 37-80.
- Fumerton, R. (1980.) Induction and Reasoning to the Best Explanation. **Philosophy of Science**, 47: 589-600.
- Geach, P. (1965) Assertion. **Philosophical Review**, 74: 449-465.
- Gibbard, A. (1990) **Wise Choices, Apt Feelings**. Cambridge, MA: Harvard University Press.
- Gibbard, A. (2003) **Thinking How to Live**. Cambridge, MA: Harvard University Press.
- Goodman, N. (1954) **Fact, Fiction and Forecast**. Cambridge, MA: Harvard University Press.
- Hempel, C. (1962) "Explanation in Science and History." In Colodny, R.C. (ed.) **Frontiers of Science and Philosophy**. Pittsburgh: The University of Pittsburgh Press. pp. 9-19.
- Hempel, C. (1965) **Aspects of Scientific Explanation and other Essays in the Philosophy of Science**. New York: Free Press.
- Hirsch, E. (1993) **Dividing Reality**. Oxford: Oxford University Press.

- Hofweber, T. (2009) "Ambitious, Yet Modest, Metaphysics". In Chalmers, D., Manley, D, and Wasserman, R. (eds.) **Metametaphysics: New Essays on the Foundations of Ontology**. Oxford: Oxford University Press. pp. 260-289.
- Hookway, C. (1988) **Quine: Language, Experience and Reality**. Oxford: Polity Press.
- Horwich, P. (1993) Gibbard's Theory of Norms. **Philosophy and Public Affairs**, 22: 67-79.
- Hudson, H. (2007) Simples and Gunk. **Philosophy Compass** 2 (2): 291-302.
- Hume, D. (1975) **Enquiries Concerning the Human Understanding and Concerning the Principles of Morals**, 3rd Edition. Oxford: Clarendon Press.
- Jenkins, C. (2013) "Explanation and Fundamentality." In Schnieder, B., Steinberg, A., and Hoeltje, M. (eds.) **Ontological Dependence, Supervenience, and Response-Dependence**. Munich: Philosophia Verlag GmbH. pp. 211-242.
- Jenkins, C. (2011) Is Metaphysical Dependence Irreflexive? **The Monist**, 94 (2): 267-76.
- Kim, J. (1994) Explanatory Knowledge and Metaphysical Dependence. **Philosophical Issues**, 5: 51-69.
- Joyce, R. (2009) "Moral Anti-Realism." In Zalta, E. (ed.) **The Stanford Encyclopedia of Philosophy** (Summer 2009 Edition). [Online] Available from: <http://plato.stanford.edu/archives/sum2009/entries/moral-anti-realism/> [Last accessed July 2014].
- Klein, P. (2005) Infinitism is the Solution to the Regress Problem. In Setup, M. and Sosa, E. **Contemporary Debates in Epistemology**. Oxford: Blackwell. pp. 131-139.
- Kriegel, U. (2013) The Epistemological Challenge of Revisionary Metaphysics. **Philosopher's Imprint**, 13 (12): 1-30.
- Ladyman, J. (2009) "Structural Realism." In Zalta, E. (ed.) **The Stanford Encyclopedia of Philosophy** (Spring 2014 Edition). [Online] Available from: <http://plato.stanford.edu/archives/sum2009/entries/structural-realism/>. [Last accessed July 2014].
- Ladyman, J. & Ross, D. (2007) **Everything Must Go: Metaphysics Naturalized**. Oxford: Oxford University Press.
- Langton, R. and Lewis, D. (1998) Defining "intrinsic". **Philosophy and Phenomenological Research**, 58 (2): 333-345.
- Lewis, D. (1970) How to Define Theoretical Terms. **Journal of Philosophy**, 67 (13): 427-446.
- Lewis, D. (1983). New Work for a Theory of Universals. **Australasian Journal of Philosophy**, 61: 343-377.
- Lewis, D. (1984). Putnam's Paradox. **Australasian Journal of Philosophy**, 62 (3): 221-36.
- Lewis, D. (1986a) **Philosophical Papers Vol. II**. Oxford: Oxford University Press. pp. 214-240.

- Lewis, D. (1986b) **On the Plurality of Worlds**. Oxford: Blackwell.
- Lewis, D. (2009) Ramseyan Humility. In Braddon-Mitchell, D., & Nola, R. (eds.) **Conceptual Analysis and Philosophical Naturalism**. Cambridge, MA: MIT Press. pp. 203-222.
- Liggins, D. (2012) "Truth-makers and dependence." In Correia, F. and Schnieder, B. (eds.) **Metaphysical Grounding: Understanding the Structure of Reality**. Cambridge: Cambridge University Press. pp. 254-271.
- Loar, B. (1990) Phenomenal States. **Philosophical Perspectives**, 4: 81-108.
- Lowe, E. J. (2010) Ontological Dependence. In Zalta, E. (ed.) **The Stanford Encyclopedia of Philosophy** (Spring 2010 Edition). [Online] Available from: <http://plato.stanford.edu/archives/spr2010/entries/dependence-ontological/> [last accessed June 2014].
- Macarthur, D. And Price, H. "Pragmatism, Quasi-Realism and the Global Challenge." In Misak, C. (ed.) **The New Pragmatists**. Oxford: Oxford University Press. pp. 91-120.
- MacBride, F. (2014) Truthmakers. In Zalta, E. (ed.) **The Stanford Encyclopedia of Philosophy** (Spring 2014 Edition). [Online] Available from: <http://plato.stanford.edu/archives/spr2014/entries/truthmakers/>. [Last accessed July 2014].
- Molnar, G. (2000) Truthmakers for Negative Truths. **Australasian Journal of Philosophy**, 78 (1): 72-86.
- Morganti, M. (2009) Ontological Priority, Fundamentality and Monism. **Dialectica**, 63 (3): 271-288.
- Mackie, J. (1977) **Ethics: Inventing Right and Wrong**. New York: Penguin.
- Olsson, E. (2005) **Against Coherence: Truth, Probability and Justification**. Oxford: Oxford University Press.
- Popper, K. (2005) **The Logic of Scientific Discovery**. [Online] Available from: <http://strangebeautiful.com/other-texts/popper-logic-scientific-discovery.pdf>. [Last accessed July 2014].
- Price, H. (2011) "Expressivism for Two Voices." In Knowles, J. and Rydenfelt, H. (eds.) **Pragmatism, Science and Naturalism**. Frankfurt am Main: Peter Lang. pp. 87-113.
- Putnam, H. (1977) Realism and Reason. **Proceedings of the American Philosophical Association**, 50: 483-498.
- Quine, W. V. O. (1948) On What There Is. **Review of Metaphysics**, 2 (5): 21-36.
- Quine, W.V.O. (1960) **Word and Object**. Cambridge, MA: MIT Press.
- Quine, W. V. O. (1966) *The Ways of Paradox* (New York: Random House).
- Quine, W. V. O. (1969) **Ontological Relativity and Other Essays**. New York: Columbia University Press.

- Quine, W. V. O. and Ullian (1970) **The Web of Belief**. New York: Random House.
- Ridge, M. (2006) Ecumenical Expressivism: Finessing Frege. **Ethics**, 116 (2). pp. 302-336
- Rodriguez-Pereyra, G. (2005) "Why Truthmakers?" In Beebe, H. and Dodd, J. (eds.) **Truthmakers: the contemporary debate**. Oxford: Oxford University Press. pp. 17-31.
- Rosen, G. (2010) "Metaphysical dependence: Grounding and reduction." In Hale, B. and Hoffman, A. (eds.) **Modality: Metaphysics, Logic, and Epistemology**. Oxford: Oxford University Press. pp. 109-136.
- Russell, B. (1927) **The Analysis of Matter**. London: Routledge Kegan Paul.
- Salmon, W. (1984) **Scientific Explanation and the Causal Structure of the World**. Princeton: Princeton University Press.
- Schaffer, J. (2003) Is There a Fundamental Level? **Noûs**, 37 (3): 498-517.
- Schaffer, J. (2004) Two Conceptions of Sparse Properties **Pacific Philosophical Quarterly**, 85 (1): 92-102.
- Schaffer, J. (2009) "On What Grounds What." In Chalmers, D., Manley, D, and Wasserman, R. (eds.) **Metametaphysics: New Essays on the Foundations of Ontology**. Oxford: Oxford University Press. pp. 347-383.
- Schaffer, J. (2010a) Monism: The Priority of the Whole. **Philosophical Review**, 119 (1): 31-76.
- Schaffer, J. (2010b) The Least Discerning and Most Promiscuous Truthmaker. **Philosophical Quarterly**, 60 (239): 307-324.
- Schaffer, J. (2010c) The Internal Relatedness of All Things. **Mind**, 119 (474): 341-376.
- Schaffer, J. (2012) "Grounding, Transitivity and Contrastivity." In Correia, F. and Schneider, B. (eds.) **Metaphysical Grounding: Understanding the Structure of Reality**. Cambridge: Cambridge University Press. pp. 122-138.
- Schaffer, J. (2014) Review of Sider's 'Writing the Book of the World'. **Philosophical Review**, 123 (1): 125-129.
- Schnieder, B. (2006) A Certain Kind of Trinity: Dependence, Substance, Explanation. **Philosophical Studies**, 129 (2): 393-419.
- Schnieder, B. (2010) Expressivism Concerning Epistemic Modals. **Philosophical Quarterly**, 60 (240): 601-615.
- Sider, T. (1993) Van Inwagen and the Possibility of Gunk. **Analysis**, 53: 285-89.
- Sider, T. (2009) "Ontological Realism." In Chalmers, D., Manley, D, and Wasserman, R. (eds.) **Metametaphysics: New Essays on the Foundations of Ontology**. Oxford: Oxford University Press. pp. 384-423
- Sider, T. (2011) **Writing the Book of the World**. Oxford: Oxford University Press.

- Simons, P. (1987) **Parts: A Study in Ontology**. Oxford: Oxford University Press.
- Skyrms, B. (1977) Resiliency, Propensity, and Causal Necessity. **Journal of Philosophy**, 74 (11): 704-713.
- Stanley, J. (2001) Hermeneutic Fictionalism. **Midwest Studies in Philosophy**, 25 (1): 36-71.
- Stoljar, D. (1993) Emotivism and Truth Conditions. **Philosophical Studies**, 70: 81-101.
- Strawson, P. F. (1959) **Individuals: An Essay in Descriptive Metaphysics**. London: Methuen.
- Taylor, B. (1993) On Natural Properties in Metaphysics. **Mind**, 102 (405): 81-100.
- Thomasson, A. **Ordinary Objects**. Oxford: Oxford University Press.
- Timmons, M. (1999) **Morality Without Foundations**. New York: Oxford University Press.
- Trogon, K. (2013a) Grounding: Necessary or Contingent? **Pacific Philosophical Quarterly**, 94: 465-485.
- Trogon, K. (2013b) An Introduction to Grounding. In Hoeltje, M., Schnieder, B., and Steinberg A. (eds.) **Varieties of Dependence**. Munich: Philosophia Verlag GmbH. pp. 97-122.
- van Fraassen, B. (1980) **The Scientific Image**. Oxford: Oxford University press.
- van Roojen, M. (2014a) "Moral Cognitivism vs. Non-Cognitivism." In Zalta, E. (ed.) **The Stanford Encyclopedia of Philosophy** (Spring 2014 Edition). [Online] Available from: <http://plato.stanford.edu/archives/sum2014/entries/moral-cognitivism/> [Last accessed July 2014].
- van Roojen, M. (2013b) "Embedding Problem Response Strategies: supplement to Moral Cognitivism vs. Non-Cognitivism." In Zalta, E. (ed.) **The Stanford Encyclopedia of Philosophy** (Spring 2014 Edition). [Online] Available from: <http://plato.stanford.edu/archives/sum2014/entries/moral-cognitivism/> [Last accessed July 2014].
- Varzi, A. (2009) "Mereology." In Zalta, E. (ed.) **The Stanford Encyclopedia of Philosophy** (Spring 2014 Edition). [Online] Available from: <http://plato.stanford.edu/archives/spr2014/entries/mereology>. [Last accessed July 2014].
- D. Wallace (2010) "Decoherence and Ontology: or, How I Learned to Stop Worrying and Love FAPP." In Saunders, S., Barrett, J., Kent, A. and Wallace, D. (eds.) **Many Worlds? Everett, Quantum Theory, and Reality**. Oxford: Oxford University Press.
- Walton, K., (1993) Metaphor and Prop Oriented Make-Believe. **European Journal of Philosophy** 1: 39-57.
- Weatherson, B. (2003) What Good are Counterexamples? **Philosophical Studies**, 115: 1-31.
- Williams, J. R. G. (2007) Eligibility and inscrutability. **Philosophical Review**, 116 (3): 361-399.

- Williams, J. R. G. (2012) Requirements on Reality. In Correia, F. and Schnieder, B. (eds.) **Metaphysical Grounding: Understanding the Structure of Reality**. Cambridge: Cambridge University Press. pp. 165-185.
- Williamson, T. (2007) **The Philosophy of Philosophy**. Oxford: Blackwell.
- Wilson, A. (manuscript) Metaphysical Causation.
- Wilson, J. (forthcoming) No Work for a Theory of Grounding. **Inquiry**.
- Woodward, J. "Scientific Explanation." In Zalta, E. (ed.) **The Stanford Encyclopedia of Philosophy** (Winter 2011 Edition). [Online] Available from: <http://plato.stanford.edu/archives/win2011/entries/scientific-explanation/> [Last Accessed July 2014].
- Worrall, J. (1989) Structural Realism: The Best of Both Worlds? **Dialectica**, 43: 99–124.
- Wright, C. (2003) **Saving the Differences: Essays on Themes from Truth and Objectivity**. Harvard: Harvard University Press.
- Yablo, S., (1998) Does Ontology Rest on a Mistake? **Proceedings of the Aristotelian Society, Supplementary Volume**, 72: 229-6.
- Yablo, S. (2001) "Go Figure: A Path Through Fictionalism." In French, P. and Wettstein, H. (eds.) **Midwest Studies in Philosophy Volume XXV: Figurative Language**. Oxford: Blackwell. pp. 72–102.
- Yablo, S. (2005) "The Myth of the Seven." In Kalderon, M. (ed.) **Fictionalism in Metaphysics**. Oxford: Clarendon. pp. 88–115.
- Yablo, S. (2014) **Aboutness**. Princeton: Princeton University Press.