

The Role of Truth in Metaphysics

Jade Emma Fletcher

Submitted in accordance with the requirements for the degree of

Doctor of Philosophy

The University of Leeds

School of Philosophy, Religion, and History of Science

July 2018

The candidate confirms that the work submitted is her own, except where work which has formed part of jointly-authored publications has been included. The contribution of the candidate and the other authors to this work has been explicitly indicated below. The candidate confirms that appropriate credit has been given within the thesis where reference has been made to the work of others.

The work in Chapter One of the thesis has appeared in publication as follows: '(Once Again) Lewis on the Analysis of Modality', in *Synthese* (2018), Divers, J. & Fletcher, J.

I was responsible for sections two and four.

The contribution of the other author was sections three and four.

This copy has been supplied on the understanding that it is copyright material and that no quotation from the thesis may be published without proper acknowledgement.

The right of Jade Emma Fletcher to be identified as Author of this work has been asserted by her in accordance with the Copyright, Designs and Patents Act 1988.

© 2018 The University of Leeds and Jade Emma Fletcher

Acknowledgements

First and foremost I want to thank my supervisor, John Divers. I don't think I can quite do justice to the role you've played over the last four years. I am so grateful for all the tireless support and critical engagement with my work. Not only is this thesis much better for it, it made me a much more careful and thoughtful scholar. Thank you also to my advisors, Heather Logue and Léa Salje. Your professional and pastoral support was invaluable. I am grateful for financial support from the University of Leeds 110 Anniversary Scholarship and the Analysis Trust Studentship, without whose backing I would have been unable to complete this work. Thank you to my examiners, Tim Button and Robbie Williams, for an invigorating *viva voce*.

I'm very grateful to have had the opportunity to be part of a number of wonderful communities over the last four years. There are many wonderful philosophers and humans in Philosophy at the University of Leeds who have contributed to my wellbeing and intellectual betterment. Special thank you to Simon Hewitt, for all coffee and helping me get my head around neo-Fregeanism, Andrew McGonigal, for first introducing me to the realism/anti-realism debate, and Robbie Williams for being an inspiring Masters supervisor. Thank you to the amazing philosophy postgraduate community at Leeds: Dani Adams, Sarah Adams, John Baldari, Sam Careelmont, Fabio Ceravolo, Marc Cole, Adina Covaci, John Fernley, David Heering, Lewis Hickley, Olof Leffler, Gary Mullen, Simon Newey, Christina Nick, Tadgh Ó Laoghaire, John Parry, Emily Paul, Mat Rees, James Rimmer, Bryan Ross, Nahuel Sznajderhaus, Nick Tasker, Alison Toop, Andrea White, and Nathan Wood. All the beer and laughs made the last four years some of the best. An extra special huge thank you to Will Gamester; the chats about truth were not only helpful and illuminating, but made it all seem like fun. Thank you to Michael Bench Capon and Thomas Brouwer for providing a source of encouragement since the first year of my undergrad. Thank you to my wonderful Leeds Girls: Laura Sellers,

Coreen McGuire, Emily Herring, and Alice Murphy. I am so grateful to have met such a wonderful group of phenomenal women, and I couldn't have done it without the cocktails and love. Thank you to Jayne at the Leeds student counselling service, who helped me navigate through the process.

I was very lucky to spend the final year of my PhD as a visiting researcher at the University of Glasgow. I'm ever so grateful to the staff and the postgraduate community for making me feel so welcome. Thank you to Ben Colburn for being so kind and encouraging, to James Humphreys, Daniel Abrahams, and Nathan Kirkwood and to the 208 office crew, Eilidh Harrison, María Pía Méndez Mateluna, Joaquim Giannotti, and Adriana Deuterium. I am grateful to the department at Glasgow for asking me to design and teach a Philosophy of Art course, and thank you to my students for keeping me going and excited about philosophy over the last hurdle. I've also been fortunate to have been able to present my work at some fantastic conferences. Being part of a supportive and exciting international community has been a real gift. My work has benefitted a great deal from many of these exchanges. I'd like to thank Katharine Jenkins, Amie Thomasson, and Sally Haslanger in particular.

I owe a huge debt of gratitude to my patient and understanding friends and family. To Sophie Noble, Saskia McCracken, and Grace Graznak: you are the best. Thank you also to my wonderful friends Alex Bovy, Nick Marsh, Richard Bellis, and Sarah Power. I owe an extra special thank you to Alex Roberts; I'm confident that I wouldn't be half the philosopher I am if I hadn't spent our entire undergraduate degree trying to keep up with you. Thank you to my wonderful mum and dad, Sian and Lindsay, this thesis wouldn't exist without you for many reasons. Thank to you my aunty Lin, your love, acceptance and support mean so much. Thank you to my aunty Mary and cousin Liam; to my in-laws, Lorraine, Peter, Gary, Sophia; and to my sister-in-law, Beth. Thank you to my grandfather, Roy Harris, to whose memory this thesis is dedicated.

Finally, thank you to my wonderful husband Sam. Go Team. W&P.

Abstract

Various areas within analytic philosophy make some central appeal to the concept of truth: the theory of meaning, metaphysics, and epistemology. They all place certain demands on the concept, and this suggests a network in which a theoretical term, ‘truth’, is implicitly and partially defined. This invites a question: is there a stable univocal concept at work – one which picks out a unique property that is capable of playing this complex role? This thesis is a constructive project which investigates this question. I aim to ascertain how different constraints placed on the concept of truth in the theory meaning can in turn constrain the role that truth can play in metaphysics. I offer six systematic case studies all of which offer different conceptions of relationship between the role assigned to truth in the theory of meaning and in metaphysics. Ultimately, I contend that paying close attention to the role of truth in metaphysics places serious pressure on the possibility of certain kinds of robust realist metaphysical projects. My thesis thus significantly contributes to the ongoing debate in metametaphysics about the status and proper method of contemporary metaphysics.

Contents

Introduction.....	11
I. The Project	11
II. The Plan	14
Chapter One: David Lewis and Philosophical Analysis	17
I. Introduction.....	17
II. Analysis, Meaning, and Truth	18
i. What is Analysis?	18
ii. How is Analysis Constrained?	23
III. Analysis and Metaphysics	28
iii. Analysis and Truthmakers	29
iv. Analysis as a Guide to Metaphysics	31
v. Semantic Analysis and Metaphysical Reduction	37
IV. Metaphysics and Truth.....	40
vi. The Role of Truth in Metaphysics	40
vii. The Metaphysical Enterprise	44
V. Conclusion	45
Chapter Two: Heather Dyke and Unmediated Metaphysics	47
I. Introduction.....	47
II. The Representational Fallacy – Overview.....	48
i. Situating the Project	48
ii. The Problematic Methodology	50
iii. The Proposed Alternative.....	52
III. The Representational Fallacy – Critique.....	54
iv. All Metaphysics as Realist Metaphysics.....	55
v. Truthmaking and the Representational Fallacy.....	58
vi. Motivating the Role of Truth.....	63
IV. Conclusion	69

Chapter Three: Ted Sider and Truth in a Metaphysical Language	71
I. Introduction.....	71
II. Sider’s project.....	72
i. Structure and Metaphysics	72
ii. The Method.....	79
III. The Status of Metaphysics.....	83
iii. The Epistemology of Structure	84
iv. Knee-Jerk Realism.....	87
v. Unavoidable Metaphysics.....	88
IV. Language Truth and Metaphysics.....	91
vi. Metaphysical and Linguistic Analysis	91
vii. The Role of Truth in Metaphysics	97
V. Conclusion	101
Chapter Four: Donald Davidson, and Semantic and Ontological	
Indeterminacy.....	103
I. Introduction.....	103
II. Truth, Meaning, and Interpretation	104
i. The Theory of Truth and the Theory of Meaning	104
ii. Radical Interpretation	109
III. Davidsonian Metaontology	112
iii. Reference and Ontology.....	112
iv. Semantic Indeterminacy and Ontology Relativity	115
IV. Truth and Metaphysics.....	121
v. Comparisons with Lewis	121
V. Conclusion	125
Chapter Five: Bob Hale, Amie Thomasson, and Truth for	
Metaontological Deflationists	127
I. Introduction.....	127
II. Hale.....	128
i. Introducing neo-Fregeanism.....	128
ii. Neo-Fregeanism Generalised	132
III. Thomasson	136
iii. Situating Easy Ontology	136

iv.	The Easy Approach and Its Consequences.....	140
IV.	Truth for Metaontological Deflationism.....	147
v.	An Important Role For Truth.....	147
vi.	Three Commitments of Metaontological Deflationism.....	150
vii.	Truth for Metaontological Deflationists.....	153
V.	Conclusion.....	156
Chapter Six: Michael Dummett, Hilary Putnam, and The		
Metaphysical Consequences of Epistemically Constrained		
Truth..... 157		
I.	Introduction.....	157
II.	Dummett.....	159
i.	Truth and Meaning.....	160
ii.	Truth and Realism.....	165
III.	Putnam.....	168
iii.	Putnam’s Attack on Realist Truth.....	170
IV.	Truth and Metaphysics.....	176
iv.	Truth and Revisionism.....	178
v.	Evaluating the Conditional.....	182
V.	Conclusion.....	188
Conclusion.....		189
Bibliography.....		195

Introduction

I. The Project

This project sits within an area of inquiry known as Metametaphysics. Metametaphysics is that area of analytic philosophy which aims to call in to question, or provide arguments for, the legitimacy or nature of metaphysics. The almost mythologised resurgence of serious metaphysics in the latter half of the twentieth century invited some questions: how should serious metaphysics be done? What constrains metaphysical inquiry? Is metaphysics even possible? If metaphysics is possible, what is the status of its theories? Are they in the business of saying something about what the world is like independent of our representation of it? Or are such theories necessarily constrained by the type of creatures doing the inquiring? This thesis is concerned with investigating a candidate constraint on metaphysical inquiry, and in so doing, providing conditional answers to some of the central question of metametaphysics.

I am interested in the role of truth in metaphysics. Why should metaphysicians care about truth? Metaphysicians can be found saying things like the following: the goal of metaphysical inquiry is to provide a true picture of what the world is like; metaphysical theories correspond to the fundamental level of reality; truth is what a representation has when it corresponds to non-representational reality; when a representation is true, it is reality that makes it so. They can also sometimes be found worrying about the status of their discipline. This worry sometimes takes the following form: perhaps metaphysics cannot tell us anything true about what the world is like. We can see, then, that talk of truth may potentially be deployed as a means of vindicating metaphysical inquiry. One of the routes to legitimacy for metaphysics is to show that the theories it generates are true, and that its methods are epistemically responsible because they are truth conducive.

There is another area of analytic philosophy that has a preoccupation with truth. The philosophy of language, in particular the theory of meaning, makes a central appeal to the concept. Philosophers of language can be found saying things like: truth is what speakers aim at when they make assertions; understanding a sentence consists in knowing the conditions under which it would be true; synonymous expressions can be substituted in a sentence without altering the sentence's truth value; the truth value of a complex expression depends upon the truth value of its component expressions. Whereas the metaphysical deployment of the concept is supposed to secure some kind of relation between representations and non-representational reality, the deployment in the philosophy of language is centrally concerned with the requirements of knowing and using a language.

Truth is thus required to play a complex theoretical role. This generates some important questions: (1) what demands do each of these enterprises place on the concept of truth, (2) how do these two deployments of the concept interact, and (3) is there a concept which can satisfy all of the demands placed upon it? Do commitments to a particular view of the role of truth on one side, force you into commitments about the role of truth on the other? That is: if we shape our concept of truth primarily with the demands of metaphysics in mind, does this put pressure on its aptitude to figure in an adequate theory of meaning? Conversely, if we delineate our concept of truth primarily with the theory of meaning in mind, does this put pressure on its aptitude to figure in metaphysics?

In this project I treat 'truth' as a theoretical term. I am trying to ascertain whether there is a consistent set of demands that can be placed on the concept. I will take it as a constraint on the acceptable deployment of the concept that it cannot make it opaque or mysterious how truth can play the role assigned to it. That is: in order to be entitled to use the concept some kind of principled story must be told about how truth gets into the picture in the first place. Such a story would show why the concept of truth is not simply dispensable in the way that redundancy theorists have advocated. This is not to say that the concept of truth needs to be explicitly defined. Indeed, and to forestall a misconception about the extent of my project, I am not in the business of

providing a general theory about what truth is. Rather, drawing on the different sets of demands placed upon the concept of truth as it is used in metaphysics and the theory of meaning, I aim to determine the theoretical and explanatory limits of that concept.

So the first major task is to investigate whether there is a consistent and explanatorily potent concept which can play a role both in the theory of meaning and in metaphysics. Having established a role for truth in metaphysics, the second major task is to investigate the conception of metaphysics with which this role for truth furnishes us. It is here where I intend to make some headway in metametaphysics. In particular, I want to determine what kinds of metaphysical projects are possible given the role assigned to truth. What I offer is a taxonomy of the theoretical space relevant to these concerns. I aim to provide a theoretical road map which details the commitments, choice points, and assumptions pertinent to different approaches to thinking about the role of truth in metaphysics. If we want to get a handle on the proper role of truth, where do we start? What does this starting place tell us or dictate about the concept of truth? What does this starting place tell us or dictate about the nature or possibility of metaphysics?

So what are the options? First, maybe there is a univocal notion of truth at work across both the theory of meaning and metaphysics, and there is a robust conception of both of these projects. If this first option is to be a contender it needs to be able to provide a story as to why should it be that correspondence to reality has anything to do with using a language. Second, maybe the concept of truth as it functions in the theory of meaning for natural language is not the notion with which metaphysicians should be concerned. Perhaps rather metaphysics requires a special metaphysical notion of truth. Finally, maybe the only well motivated notion of truth that metaphysics can entertain must be the one that functions in the theory of meaning. The question then becomes: what are the metaphysical consequences of this? Given the constructive nature of my project, I do not consider those positions according to which truth is completely dispensable from either semantic or metaphysical inquiry.

This thesis proceeds by offering six systematic case studies. Investigations into metaphysical methodology are by their nature highly speculative. It is,

therefore, even more important that steps are taken to prevent such inquiry from becoming hopelessly stratospheric. Approaching questions of metaphysical methodology through case studies which exemplify the different theoretical options has the virtue of constraining inquiry by rooting it in the details of those worked-out philosophical positions. Of course, any findings will be of limited interest if the lessons learned from reflecting on these case studies cannot be generalised. My present investigation does not, however, have this flaw. I take it that each case study I evaluate exemplifies different common themes and, whilst my evaluations of each proposal may not uniformly transfer over to similar positions, they do generate an instructive set of concerns with which any similar position should contend.

II. The Plan

I offer the following six case studies as a means of exploring and evaluating the available theoretical options:

In Chapter One I evaluate David Lewis's project of philosophical analysis. I take Lewis's project as my starting point as it promises robust conceptions of both meaning and metaphysics, with a central and integrated role for truth in both of these inquiries. In Lewis's system, truth both anchors his theory of meaning by playing an indispensable and central role in the semantic component of that theory, and it supervenes on being. I present Lewis's project of philosophical analysis as the method by which each inquiry is conducted, as well as how they connect and support each other. My central claim about Lewis and the role of truth in metaphysics is that the notion of truth does not have work to do beyond its role in semantic analysis and a broad supervenience claim. Thus, in order to sustain a robustly realist metaphysical enterprise, Lewis must utilise distinct metaphysical methods.

In Chapter Two I evaluate Heather Dyke's philosophical project in *Metaphysics and the Representational Fallacy*. Dyke offers a metaphysical conception of truth and her project is representative of a certain animosity towards taking meaning

and language to be in some sense authoritative when doing metaphysics. This chapter therefore takes Dyke's position as a case study for the prospects of an unmediated approach to metaphysics. I argue both that Dyke's arguments against taking language seriously when doing metaphysics, and that her proposed alternative methodological strategy are unsuccessful. I argue that if we propose to make use of the concept of truth in a metaphysical framework that distances itself from the role assigned to truth in the theory of meaning, then we must have a story to tell about how that concept can play the role assigned to it.

In Chapter Three I evaluate a language-first approach to metaphysics, where the language in question is not natural language but rather a special metaphysical language. I take as my primary case study Ted Sider's project in *Writing the Book of the World* (2011). Sider's project offers a promising synthesis between approaches to metaphysics which take meaning and truth to be altogether orthogonal to metaphysical inquiry, and those which take metaphysics to be in some sense posterior to considerations about language. The method of metaphysics is therefore centrally concerned with meaning and truth, but not as these notions function in the theory of meaning for natural language. I argue that for Sider's approach to be viable he must articulate an adequate theoretical role for truth, but that as it stands his approach is lacking such an articulation.

In Chapter Four I evaluate Donald Davidson's conception of the role of truth in metaphysics, and the status of the subsequent metaphysical inquiry. Davidson's project presents an interesting case study. A semantic conception of truth is his central theoretical notion, which earns its keep through the role it plays in the theory of meaning. Metaphysical commitments are then generated when the true sentences of the language are analysed to reveal quantificational structure. I show how Davidson's project illuminates how the status of metaphysical inquiry can be transformed depending on claims made about other semantic concepts. I argue that if you think that a truth conditional analysis of language is a prerequisite for generating metaphysical commitments, but that the interpretation of that language is indeterminate, then these

commitments together limit the type of metaphysical project in which you can engage.

In Chapter Five I evaluate the prospects for a deflationary approach to the project of ontology. My discussion focuses on the neo-Fregeanism of Bob Hale and the neo-Carnapian position championed by Amie Thomasson. Such approaches advocate for a view of metaphysics according to which ontology is exhausted by semantically analysing the true sentences of the language. If viable, these approaches, which I call *metaontological deflationism*, promise a realist ontology with a transparent epistemology. I argue that for such an approach to be viable, for it to be able to deliver on its promises, it must articulate an adequate conception of truth. I contend that as it stands the approach is lacking such an articulation. My aims in the chapter are to explain why the metaontological deflationist must articulate an adequate conception of truth; argue that there are three constraints on how their operative conception of truth should be understood; and present a positive proposal for how the metaontological deflationist should approach the articulation of their conception.

In Chapter Six I evaluate how commitment to an epistemically constrained conception of truth can affect the possibility and nature of metaphysical inquiry. I consider two historically influential arguments for the claim that all truths are knowable: one from Dummett and one from Putnam. According to this approach, doing justice to the role of truth in language requires rejecting both an epistemically unconstrained conception of truth and realist metaphysics. The aim of this chapter is not to establish whether or not truth is in fact epistemically constrained. The aim of this chapter is rather to evaluate the supposed connection between truth and realism. I defend the following conditional: if truth is epistemically constrained, then this has unavoidable metametaphysical consequences. Defending this conditional requires some significant revision of what Dummett and Putnam take to be the realist/anti-realist consequences of their respective arguments.

Chapter One: David Lewis and Philosophical Analysis

I. Introduction

The purpose of this chapter is to evaluate David Lewis's conception of the role of truth in metaphysics. I take Lewis's project as my starting point as it promises robust conceptions of both meaning and metaphysics, with a central and integrated role for truth in both of these inquiries. In Lewis's system, truth is both indispensable to semantics, and anchors his theory of meaning, and it supervenes on being, and as such is connected to his metaphysical theorising. I present Lewis's project of philosophical analysis as the method by which each inquiry should be conducted, as well as how they connect and support each other. Lewis does not provide a single, explicit, comprehensive, and authoritative statement of what he takes analysis to be. However, in many places he offers statements which partially illuminate what he takes to be constitutive of an analysis. Therefore, by taking his work as a whole, a picture of his analytical method can be reconstructed and his conception of the role of truth in metaphysics can be understood.

The chapter proceeds as follows. First, I establish the primary purpose of analysis. (§II) I argue that Lewis's brand of philosophical analysis is intrinsically and primarily concerned with giving the meanings of expressions in terms of their truth conditions. Second, I detail the relationship between analysis and metaphysics. (§III) I argue that whilst Lewis's method proposes a systematic approach to semantic and metaphysical theorising, these projects must be viewed as theoretically distinct. I propose an understanding of Lewisian analysis whereby it is indispensable to metaphysics yet a different theoretical enterprise. Finally, I consider whether the role assigned to truth in Lewis's framework constrains the nature of the resultant metaphysical project. (§IV) I

argue that the theoretical role assigned to truth in metaphysics has a limit. Consequently, Lewis must employ a metaphysical method independent of the role of truth if he is to engage in his robustly realist metaphysical enterprise.

II. Analysis, Meaning, and Truth

i. What is Analysis?

The first point to make is the frequency of Lewis's use of the phrase 'semantic analysis'.¹ We must be careful to distinguish two readings of this. There is one reading according to which 'semantic analysis' is but a subspecies of the broader category of analysis. There is a second reading according to which analysis is exhausted by 'semantic analysis'. My thesis is that it is the latter of these readings which best characterises Lewis's analytical projects: for Lewis, semantic analysis and analysis are the same thing. Much of what follows defends that this is the only plausible interpretation of *Lewisian analysis*.

To understand the aim of semantic analysis, we need to appreciate Lewis's view of language, and the relationship between semantics and truth. Lewis's conception of natural language is as a conventional practice governed by norms of assertability which allow successful interpretation and communication between truthful and trusting partners. In giving a theory of meaning Lewis wants to provide a two tier account. (Lewis 1970b, 190) One project consists in providing an abstract description of the workings of possible languages. This gives us the idea of languages as semantic and syntactic systems, where items of a language get to be associated with things in the world. There is another project which is the project of explaining the psychological, sociological, and worldly facts that underpin how particular

¹ For example see Lewis 1970a, 10; 1973b, 118; 1974b, 260; 1979a, 154; 1979b, 38; 1981, 415; 1983b, 17; 1994b, 319.

populations use one of these abstract semantic systems. Although the projects can be drawn apart, the former project is importantly constrained by the latter.

The task is thus to give a semantic theory for natural language in such a way as to be appropriately constrained by considerations to do with language use, translation, and interpretation. Lewis's primary motivation in offering the semantic theory that he does is that it should contribute to the explanation of successful communication between interlocutors. (Lewis 1976a, 139-140)

Truth then gets into the picture by playing an explanatory role in the theory of communication: it is that feature which we take sentences to have when we assert or believe them. Consequently, Lewis thinks that the best grasp we have on meaning is through giving truth conditions; he holds that "semantics with no treatment of truth conditions is not semantics". (Lewis 1970b, 190) We can see why this is so. If I utter a sentence with assertoric force then I am claiming that things are as the sentence says they are. So, for me to understand a sentence (i.e. know its meaning), I need to know how things would be if what was asserted by the sentence were true. That is: I need to know its truth conditions.

Whilst Lewis does not explicitly express his views on the function of the concept of truth in this way, I think it is a justified interpretive step for the following two reasons. First, although Lewis doesn't offer an explicit, comprehensive, and authoritative statement of how he is thinking about the role of truth, taking his work as a whole we can extract the story above. An example of such a partial statement is the following:

"We might reasonably take it as the goal of semantics to specify our prevailing rules of assertability. Most of the time, to be sure, that can best be done by giving truth conditions plus the general rule that speakers should try to be truthful, or in other words that assertability goes by probability of truth." (Lewis 1976a, 139-140)

Second, taking Lewis to think about the role of truth in the way I suggest is in keeping with his overarching methodology. We start by treating truth as a theoretical term, and then ask what role we need it to play. Thinking about truth in this way is analogous to how Lewis asks us to think about meaning.

Although Lewis does give us something which he considers to be the meanings of expressions, he doesn't take himself to have given the definitive answer on what should be construed as 'meaning'. Much like his other analyses, I take his view of meaning to be providing an account of how we should construe meaning, given the role we want it to play in our theorising. (Lewis 1986b, 49) Truth, then, first and foremost gets into the picture via the role we need it to play in the theory of communication. From here, we then construe semantics as the study of the meaning and truth conditions of whole sentences, and the contribution made by their constituent expressions to their meaning and truth conditions.

A *semantic* analysis for a particular term or expression gives the meaning for that term or expression. Lewis thought a good candidate for the meaning of sentences, names, common nouns, etc. are functions that give as a value an extension and take as argument various contextual factors on which that extension may depend. (Lewis 1970b, 194) For example, in 'Languages and Language' he claims that the meaning of a sentence is "something which, when combined with factual information about the world – or factual information about *any* possible world – yields a truth value" (Lewis 1975, 163). Likewise, in *Convention* he tells us that the meaning of sentence (insofar as this can be given without the meanings of the parts) is the function "whereby its set of interpretations depends on features of occasions of utterance" (Lewis 1969, 165). The function, therefore, provides the analysis of the analysandum and, in so doing, gives its meaning.

An analysis gives the meaning of the analysed expression. The best handle we have on the meaning of an expression is via its truth conditions. So, an analysis must provide truth conditions. The type of truth conditions that Lewis thinks should be given by analysis are *world invariant* truth conditions. World invariant truth conditions provide a uniform type of condition that applies at each world, and a sentence is true at a world if and only if it satisfies the condition at that world. It is because such truth conditions are supplied that semantic analysis can enrich our understanding of the subject matter in question. For example, Lewis is in a position to defuse certain criticisms of identity theory

because of his analysis of expressions of necessity when they occur de re.
(Lewis 1971, 49)

For Lewis, analysing a term and defining that term are very often the same thing. When discussing Armstrong's 'compulsory question'², Lewis claims Armstrong is working with an *unfamiliar* notion of analysis. (Lewis 1992, 202-203) By contrast, the 'familiar notion' is one whereby analysis is primarily a search for definitions. I take it that this is Lewis's conception of analysis. It is by supplying definitions that Lewis gives the meaning and the truth conditions for the concepts in which he is interested. In addition to distancing Lewis from Armstrong's conception of analysis (which is a search for truthmakers), we must also distance the Lewisian project of definition from the project of *Real Definition*: Lewis is interested in the definition of linguistic expressions, unlike the project of Real Definition which aims to define (in general) non-linguistic phenomena.

In 'New Work for a Theory of Universals', when talking about the analysis of duplication, Lewis claims that a particular "'definition' is no analysis."³ (Lewis 1983b, 27) The reason it is not an analysis is that it does not provide truth conditions which are world invariant. Definitions may be reference fixing, but if they only fix reference at this world, they do not provide the sort of definitions that Lewis thinks are given by an analysis. Such sentiments are reiterated in 'Causation' when he claims that much needs doing to turn a candidate definition into a defensible analysis.⁴ (Lewis 1973a, 160) The commitment motivating this claim is that an analysis must not just serve to successfully characterise something at this world, but must apply to all worlds. To further emphasize this point, in a number of places Lewis claims that although a candidate analysis must successfully systematise opinion, if it fails to provide a definition that specifies world invariant truth conditions for the

² That is: the question of unanalysed predication. See Armstrong (1980).

³ Here, Lewis is criticising analyses which propose to analyse duplication as *alikehood in arrangement of particles*.

⁴ In this case Lewis claims that regularity analyses of causation fail as they confuse causation itself with other causal relations.

analysandum, it cannot be seen as an acceptable analysis. (Lewis 1997a, 140; 1983d, 158)

Although an adequate definition must always provide world invariant truth conditions, it is worth noting that the form of the definition can differ depending on the nature of the definiendum. In his analyses of contingent concepts, such as mental state terms or causation, Lewis provides functional definitions. For example, as a matter of analytic necessity counterfactual dependence of events on other events provides the necessary and sufficient conditions for causation. (Lewis 2004, 287) This analysis holds across all worlds in virtue of its articulation in terms of world invariant truth conditions. This, however, allows that as a matter of contingent fact the thing at the actual world “on which the truth values of causal statements supervene” is the ‘biff’ kind of causation. (Lewis 2004, 287) Similarly, a mental state term such as ‘pain’ is analysed in terms of its causal role. It can then be contingently identified with a physical role-player at the actual world, and we can thus contingently derive the truth of identity theory. The idea driving this treatment is that semantically contingent expressions are those whose extensions vary from world to world. By contrast, in his analysis of a non-contingent concept, such as necessity, Lewis gives an explicit definition. Unlike the contingent cases, the truth values of necessity statements do not vary from world to world: it is the whole of logical space which determines their truth value.

Lewis provides us with a number of features that these definitions must have. They should be general, finitely specifiable, a priori, uniform, and simple. (Lewis 1997a, 144; 1979b, 41; 1974a, 115) A good illustration of Lewis’s commitment to finite stateability comes in his treatment of supervenience in ‘A Subjectivist’s Guide to Objective Chance’. When discussing supervenience, he says that we shouldn’t take it to be a doctrine of analysability, since the *metaphysical* thesis of supervenience, cannot be stated in a finite way.⁵ (Lewis

⁵ By this I do not mean that the *doctrine* of metaphysical supervenience cannot be stated finitely (e.g. no change in A-facts without a change in B-facts). Rather I mean that a full specification of the A-facts, on which the B-facts supervene, cannot be finitely stated, and thus metaphysical supervenience cannot be doctrine of analysability. This point will become clearer in section III.

1980, 111) That is: because it falls short of finite definability it cannot be part of analysis itself. (Lewis 1986b, 222)

Up to this point, it seems like Lewis is engaged in a very traditional philosophical project. Lewis thinks that starting point of systematic philosophy is semantic analysis. (Lewis 1983a, 112, 1983b, 20-21) However, Lewis's conception of analysis is not a mere restatement of a more traditional conception. For example, in 'Reduction of Mind', he tells us:

“Arbiters of fashion proclaim that analysis is out of date. Yet without it, I see no possible way to establish that any feature of the world does or does not deserve a name drawn from our traditional mental vocabulary. We should repudiate not analysis itself, but only some simplistic goals for it. We should allow for semantic indecision: any interesting analysandum is likely to turn out vague and ambiguous. Often the best that any one analysis can do is to fall safely within the range of indecision. And it should allow for semantic satisficing: analysis may reveal what it would take to deserve a name perfectly, but imperfect deservers of the name may yet deserve it well enough.”

(Lewis 1994b, 298)

Thus, whilst Lewis's brand of philosophical analysis is intrinsically and primarily a matter of giving the meanings for expressions, in terms of their truth conditions, analysis need not rule out semantic indecision, and an analysis can still be useful even when it reveals that there is nothing that satisfies it perfectly.

ii. How is Analysis Constrained?

For Lewis the aim of semantic analysis is to explain why the things that we take to be true can be so taken. In determining belief contents, through the process of radical interpretation, we should be charitable to agents, and thus assign true

beliefs where possible.⁶ The assignment of truth conditions to the sentences we hold to be true is, in turn, constrained by this commitment to charity. Thus, if we seem to pre-theoretically say true things about the analysandum, Lewis is going to want to offer an analysis, and thus truth conditions, which respect this feature of our discourse. This explains Lewis's general rejection of error-theoretic positions. For example, Lewis claims that an "analysis that imputes widespread error is *prima facie* implausible." (Lewis 1986d, 175) In this section I want to say more about how common opinion acts as the data for and a constraint on analysis. I take this to be illuminating as it helpfully draws out a number of features of analysis, which helps to set up some of the issues I discuss in section III.

Lewis tells us that our job as theorists is to systematize the mass of common opinion into some kind of coherent picture: "A metaphysician's analysis of mind is an attempt at systematizing our opinions about mind." (Lewis 1973b, 88) Whilst merely systematizing common opinion is not *sufficient* for providing an adequate analysis, in many places Lewis presents it as a necessary feature of a satisfactory theory. (Lewis 1974a, 111-114) For Lewis, 'common opinion' refers to those beliefs or statements which we collectively take to be true as language users. This will include, in addition to folk common sense, some well-established and widely held scientific beliefs. For example, among our analysis-relevant pre-philosophical opinions are the beliefs that there are more particles in my coffee mug than there are dinosaurs, and that volcanoes are caused by tectonic activity. Common opinion, therefore, must be pre-philosophical opinion, but may still include some well-established theoretical scientific opinions.

Lewis takes it as non-negotiable that analysis must (to some extent) accommodate the platitudes of common sense. Whilst the analysis we end up with does not have to have all those platitudes as direct components, it must still be compatible with them. To say that common opinion is the data for analysis does not necessarily mean that it is the common opinion itself that we

⁶This isn't to say that the principle of charity is the only constraint on interpretation. See (1974a) for Lewis's discussion of how charity fits into the wider rationalisation story.

are analysing: common opinion acts, in many cases, as *the signifier of the explanandum*. When we try to construct a theory of, say, modality, it is possibility and necessity in which we are interested. However, our best guide to that subject matter is to be given by the everyday pre-theoretic thought pertaining to the area in question. Lewis can talk both of systematizing common opinion and analysing the concept; this makes sense as it is the common opinion which helps us pick out the thing in which we are interested.

Common opinion serves another function in Lewis's analytic projects: in addition to acting as the data for analysis, common opinion also acts to constrain what constitutes an acceptable analysis. To explain the role of common opinion, therefore, it is not sufficient to just acknowledge that it fixes the relevant notion, we must explain how it further guides the analysis. In many places throughout his work Lewis tells us that if presented with two analyses, two posits, two answers to particular problems, we should prefer "the platitude of common sense to the interesting philosophical thesis." (Lewis 1976b, 56) He argues that in cases where we have a firm and uncontroversial opinion, "theory had better agree." (Lewis 1986d, 194) The guiding idea is that the plausibility of any particular analysis is seriously decreased if it diverges too far from what we ordinarily think about such matters. Thus, analysis must be conservative of common opinion. Lewis claims that to deny common opinion this role is to reject much of what we would recognise as analytic philosophical method. (Lewis 1994b, 311-312)

I think it is useful to pause briefly to consider how common opinion plays a role in connecting two of Lewis's semantic projects. I said before that Lewis's semantic theory must be appropriately constrained by considerations to do with language use, interpretation, and communication. The metasemantic story here is Lewis's project of radical interpretation, which is his foundational account of the semantic content of beliefs and sentences. According to this project, the content of beliefs is whatever it has to be in order to rationalise an agent's actions, given their evidence.⁷ From here we can work out what

⁷ I discuss Lewis's radical interpretation project at greater length in Chapter Four when contrasting it with Davidson's.

speakers' sentences mean. Radical interpretation provides an account of how it is that we should assign meaning and thus truth conditions to speakers' sentences. How does this fit together with the analytical project as I've described it? Both are in the business of giving the meaning of expressions of natural language, but they seem like different projects. Lewis does not offer an explicit comment on exactly how these projects relate. Nonetheless, I think there is a fairly natural way of fitting them together, which further highlights to role of common opinion.

Radical interpretation is the means by which we determine the content of an agent's beliefs and sentences. Semantic analysis is the means by which we give definitions of the concepts in which we are interested. The data required for semantic analysis are the commonly held beliefs of a linguistic community. Thus it is the output from the radical interpretation project which provides part of the input for the analytical project. Consider an example. We tell a rationalisation story to arrive at the claim that Stirling believes Kane is in pain. Further, Stirling believes that Kane is likely to cry, is recoiling because he's in pain, etc. Imagine all the other rationalisation stories which could be told to arrive at similar belief contents. Now, imagine we want to semantically analyse the mental state term 'pain'. To begin, we need to elicit and systematise all of these pre-philosophical common beliefs into the folk theory of pain. And from here we proceed as I outlined above.

In establishing this robust role for common opinion, we must acknowledge that Lewis treats it as defeasible: he thinks that it is permissible to deviate from common opinion. For example, he tells us that "canons of reasonable belief need not be counsels on perfection." (Lewis 1980, 85) The question is, then, under what circumstances should common opinion be respected, and when should it be overridden? Lewis does not give a categorical answer to this question. However, I think his view on the limits of common opinion as a decisive theoretical tool can be reconstructed. There are three reasons why we might not want to take common opinion to be a significant constraint on our theorising. I take them in turn.

First, our opinions about certain cases might be inconsistent. That is: we might have competing opinions about some cases. Remember, part of the job of

analysis is to systematise our opinions about the subject matter. However, we cannot systematise what is inconsistent. I take it that a good example of a case such as this is his treatment of *actuality*. When faced with a tension between our opinions about actuality, Lewis opts to side with common opinion about what sorts of things are actual, how many things are actual, and how those things relate to us, which forces him to reject common opinion that everything is actual. (Lewis 1983c, 25; 1986b, 100) The choice to go one way rather than another in this case is motivated by his wider theoretical commitments.

Second, Lewis does not think that we need to respect opinion that is not really *common* opinion. I take it that there are a number of ways this could be the case. Perhaps the opinion is “offhand”, in which case he does not think we need to adhere to it. (Lewis 1973b, 80; 1986b, 87) For example, he does not want to endorse off-hand opinion about colour metaphysics. (Lewis 1997b, 332) He also does not think we need to pay attention to opinion about extreme cases. For example, he claims that in cases of “extreme oddities” our opinions cannot be taken as a guide, as in the truth conditions of farfetched counterfactuals. (Lewis 1973b, 28) He counsels that in cases where it appears that we do have opinions about farfetched matters, we should not trust them. Thus, he rejects “naïve opinion”. (Lewis 1986b, 246-247) In addition, we also do not need respect opinion that has become saturated with theory. (Lewis 1974a, 111-112) I take it that what unites these ways in which common opinion can be safely ignored is that they are all cases in which our opinions are too far removed from the kind of pre-philosophical thought that Lewis thinks can act as a constraint on analysis.

Finally, we may diverge from common opinion because it leads to an unfavourable theory. For example, if an analysis forces us to accept something mysterious, we do not need to adhere to such opinions, as in his discussion of the counterfactual analysis of causation. (Lewis 1986d, 182) A particularly interesting manifestation of this kind of motivation to move away from common opinion is that if common opinion forces us to accept an analysis that, in turn, dictates acceptance of certain empirical theories, we need not adhere to that opinion. This point invites discussion. Much in the same way that common opinion acts as a defeasible restriction on what constitutes an

acceptable analysis, Lewis also takes scientific theory to place a restriction on his analytical projects. Lewis does not think scientific posits should be *constitutive* of an analysis. The reason for this is that Lewis does not think that scientific relationships deserve to be built into our analyses, as we could not establish such connections a priori. (Lewis 1992, 209; 1986c, 126) Two things follow from Lewis's commitment to the a prioricity of analysis. First, analysis itself should not posit empirical relations. Second, we should reject an analysis if it forces either the acceptance or the dismissal of particular scientific theories.

This latter component has two dimensions. On the one hand, Lewis does not want any particular analysis to rule out any physical theories as impossible a priori. (Lewis 1979b, 41) Although we may have good empirical reasons for rejecting a certain scientific theory, Lewis thinks that the falsity of scientific hypotheses should not be established by means of analysis. From the other direction, he does not want analysis to force the acceptance of particular scientific theories. This would be to a priori commit ourselves to things that should be established empirically. Lewis thinks that there are certain kinds of things (for example for empirical posits of natural science) which should not be determined by philosophy alone, and certainly not by analysis. (Lewis 1992, 198) Essentially, he does not want to "impose a priori answers on questions that ought to be empirical." (Lewis 1986e, 63)

III. Analysis and Metaphysics

In this section I outline Lewis's conception of the relationship between analysis and metaphysics. Having set up in the previous section that the primary purpose of analysis is giving the meanings of expressions in terms of truth conditions, I consider how we should think about this project as relating to metaphysical inquiry. I consider three hypotheses: (1) analysis should provide truthmakers, (2) analysis and metaphysical inquiry are the very same thing, (3)

analysis should provide metaphysical reductions. I argue that all of these views are misguided as interpretations of Lewis.

iii. Analysis and Truthmakers

According to the view of analysis I defended in the previous section, analysis must give truth conditions. This seems fairly innocuous. However, you might think that this claim is radically underspecified. Perhaps we should understand the request for truth conditions as a request for truthmakers. Truthmakers place a certain kind of *worldly constraint* on the satisfaction of the truth conditions: the truthmaker must exist. In the context of his discussion of Humean supervenience Lewis endorses Bigelow's claim that truth supervenes on being. (Lewis 1994a, 225) So, with respect to a particular world the totality of truths supervene on being. The supervenience claim must concern a particular part of logical space, and it must be expressed holistically. It must concern a part of logical space as, for Lewis, the totality of being is non-contingent, so the totality of truths across all of possible space cannot vary. It must be holistic as it would not be instructive, or perhaps possible, to go sentence by sentence.

However, whilst Lewis is prepared to endorse some notion of *truthmaking*, that does not commit him to an endorsement of *truthmakers*. That is: he can accept that there are some worldly conditions on truth, without a more specific commitment to truth always obtaining in virtue of the existence of an *entity*. This doesn't rule out the possibility that sometimes there might be an entity that makes a particular statement true. For example, the sentence 'Jade exists' is plausibly made true by me. However, we shouldn't take the fact that some statements have truthmakers to reveal a requirement of *truth*, but rather it shows something interesting about those particular existential statements. In resisting the request for truthmakers then, I take the thought to be that even if there are some worldly conditions on truth, we need not consequently engage in a practice of positing theoretical entities (e.g. states of affairs, Tractatian facts, etc.) in virtue of which truth invariably obtains.

The idea that Lewisian analysis should provide truthmakers is not a thesis that suggests itself from Lewis's corpus. However, he does briefly discuss the idea that analysis is primarily a quest for truthmakers. For example, when discussing Armstrong on analysis he writes:

“I suggest that Armstrong has an unfamiliar notion of analysis. Analysis is not, primarily, a quest for definitions. Rather, it is a quest for truthmakers. The ‘harlot’s privilege’ is not the privilege of using undefined terms. It is the privilege of truth without benefit of truthmakers.” (Lewis 1992, 202-203)

In identifying this ‘quest for truthmakers’ as an ‘unfamiliar notion of analysis’ I take it Lewis does not identify such a notion as his own. To add further credence to the thought that this is not what Lewis has in mind, Lewis explicitly does not commit himself to the view that all truths must have truthmakers (Lewis 1992, 204), whereas he does think that all truths must have truth conditions. This is sufficient to establish that he does not think that truth conditions should be understood in terms of truthmakers.

However, the important idea driving this commitment is that the specification of truthmakers cannot be constitutive of analysis, as the constraints on analysis require that the connections posited by the analysis must be *a priori*. Lewis writes:

“We don’t know any recipe for a talking donkey in terms of elements of the basic structure; and if we did, we still wouldn’t know it *a priori* and so wouldn’t be entitled to build it into our analyses.” (Lewis 1992, 209)

The world which supplies the truthmakers may be complex and infinite, and avails itself to empirical *a posteriori* investigation. However, we have already established that Lewis takes analysis to be specifying finite and stateable definitions. The analysis, therefore, must meet conditions which the worldly truthmakers need not. Inasmuch as this is the case, we may discharge the idea that analysis should provide truthmakers as being inconsistent both with explicit claims that Lewis makes about the ‘unfamiliar notion of analysis’ and with the implicit consequences of his other analytical commitments.

iv. Analysis as a Guide to Metaphysics

Lewis claims that analysis serves the cause of metaphysics by systematising our pre-theoretic opinions pertaining to the area in question. He tells us:

“One comes to philosophy already endowed with a stock of opinions. It is not the business of philosophy either to undermine or to justify these preexisting opinions, to any great extent, but only to try to discover ways of expanding them into an orderly system. A metaphysician’s analysis of mind is an attempt at systematizing our opinions about mind. It succeeds to the extent that (1) it is systematic, and (2) it respects those of our pre-theoretic opinions to which we are firmly attached. Insofar as it does both better than any alternative we have thought of, we give it credence... So it is throughout metaphysics.” (Lewis 1973b, 88)

There is a multiplicity of words or concepts that play certain roles in our ordinary and everyday thinking and in our various philosophical and scientific theories. (Lewis 1986b, 55) Answering the metaphysical question of what those things are turns on the question of which entities are best suited to play the everyday and theoretical roles of the concept in question. (Lewis 1986b, 184)

Taking functional analyses as a case study, consider this possible interpretation of Lewis’s view on the connection between analysis and metaphysics: analysis just is metaphysics. It is analysis that systematises our pre-theoretic thought. This systemisation provides us with a definition in terms of the relation the definiendum bears to our other concepts and beliefs. (Lewis 1972, 255) We might then think that whatever thing in the world satisfies the functional definition is part of our metaphysical theory. So, for example, Lewis tells us that: “A feature of Reality deserves the name of chance to the extent that it occupies the definitive role of chance.” (Lewis 1994a, 246) According to such a view, then, we are to construct our metaphysical theories by offering an account of what the relevant terms mean, in terms of their truth conditions. This would give truth a central and indispensable role in metaphysics.

There are a number of reasons we might think that analysis just is metaphysics. I take them in turn. First, in his discussion of the definition of theoretical terms, Lewis seems to suggest that analysis is sufficient for determining ontological commitments. For example:

“The defining of theoretical terms serves the cause of scientific realism. A term correctly defined by means of other terms that admittedly have sense and denotation can scarcely be regarded as a mere bead on a formal abacus. If it purports to name something, then if the theory that introduced it is true it does name something.” (Lewis 1970c, 79)

And,

“Because we understand the O-terms, and we can define the T-terms from them, theories are fully meaningful; we have reason to think a good theory true; and if a theory is true, then whatever exists according to the theory really does exist.” (Lewis 1972, 256)

Certainly, if we proceed as suggested, our analyses will anchor our metaphysical theories onto our subject matter and our metaphysical posits will be constrained by our analyses. Perhaps, then, we should take analysis to settle metaphysical questions? I do not think this view of Lewisian metaphysics can be right. Our ability to reliably deploy our concepts does not necessarily provide a satisfactory understanding of what sort of entities they refer to (1986b:189). Even after we have constructed the definition from the pre-theoretic thought, and even if that term appears in a true statement, a substantial amount of work still needs to be done to determine what thing in the world satisfies it.

Second, another point at which we might think that analysis gives us metaphysical commitments, is in Lewis’s discussion of semantic values. Lewis can often be found to be discussing the ‘entities’ to which the analysis commits us.⁸ However, I do not take this to be a commitment to substantive metaphysical theory. It is important when thinking about Lewis’s view of the analysis-metaphysical connection, therefore, to disambiguate between the

⁸ For some good examples see Lewis (1969), (1970b), (1970c), and (1972).

commitments generated by the semantic theory versus a more full blooded type of metaphysical commitment. When Lewis says: “If we want to give the connective an entity to be its meaning, the operation can serve this purpose” (Lewis 1973b, 47), he is providing a referent for the connective. Likewise for the case of properties: a set acts as the referent for the predicate. However, in these cases Lewis’s talk of entities should not be seen as telling us anything significant about non-semantic reality.⁹ The thought driving the distinction here is that although analysis can existentially commit us, it does so in a metaphysically open way. So, analysis can (assuming the truth of the relevant statements) commit us to the existence of something that plays the role of properties. What analysis does not commit us to is what the particular something is that plays that role. Thus, analysis (and truth) can metaphysically commit us to certain roles, but there is a further metaphysical question about what the role players are.

Third, in setting up analysis as his primary philosophical project, Lewis appears to locate it as a metaphysical option. When attempting to offer a theory of some subject matter, such as causation, colour, mind, or modality, Lewis claims that there are three ways to give an account. (Lewis 1983a, 112, 1983b, 20-21) First, we can decline the invitation and go eliminativist. This option requires that we deny the existence of the things in question, and we must then explain away any apparent commitment to them. Lewis claims this approach is generally unfavourable. Second, we can accept it as primitive. Thus we embrace the existence of the thing in question, but we are not required to offer a more basic account of it in other terms. Lewis thinks that all theories must have primitives, but prefers analyses where possible. Finally, we can offer an analysis. Eliminativism and primitivism seems to be metaphysical positions: either denying existence, or claiming that that something is metaphysically basic. In pitting the quest for analyses against these other two theoretical options, we might be tempted to think that analysis is itself a metaphysical enterprise. I think this would be a mistake. The question of analysis is rather a

⁹ By way of example, consider an error theory about colour. We can agree about what the relevant colour terms mean (i.e. we can agree on their semantic analysis), but substantially disagree once we move below the level of semantic reality.

question of methodological priority. Prior to analysis, we might think that we need to be ontologically committed to, for example, irreducible mental states. We can then engage in questions about whether we should deny their existence or take it as metaphysically basic. However, once we have analysed our mental state terms, we can see that we needn't engage this particular question at all. Rather, we need to find satisfiers for mental states' functional role. In this way, analysis may help to direct metaphysical inquiry, but isn't that inquiry itself.

Whilst it is fair to claim, therefore, that the analyses Lewis offers do *constrain* or *direct* his metaphysical inquiries, the analyses cannot be seen as *constituting* the metaphysical theories themselves. The conception of the analysis-metaphysics connection endorsed by Lewis that I want to urge here is twofold. First, there is an analytic element of metaphysical inquiry, which specifies the meaning of terms and acts to stop the metaphysical posits from floating free from what the terms mean and how they are used. Second, that metaphysical method is still, in some way yet to be unspecified, distinct from the analysis.

I take one of the primary motivations for the claim that analysis and metaphysical inquiry are not the same thing is that in a number of places Lewis tells us that analysis will only get you so far. Although analysis may determine part of the content of our metaphysical theories (e.g. the functional role of the thing in question) we do not get a full picture of the commitments. I take it that this is something that becomes pronounced when you view Lewis's corpus as a whole. Although in the early work, when his attention is directed at more logical and semantic concerns, he provides analyses of certain kinds of things, for example possibility and necessity in terms of quantification over worlds (Lewis 1968, 1970a), he leaves much of the metaphysics largely untouched. For example, he claims:

“Up to a point it makes little difference whether you believe as I do in a multitude of concrete worlds of which ours is one, or whether instead you believe as Quine does in a multitude of abstract ersatz worlds, of which one is special in that it represents that one and only concrete world. Most analyses involving possible worlds go through equally well either way.” (Lewis 1979a, 148)

Lewis thinks that there will be much of metaphysical interest left unsaid even after we have given the analysis.

In the secondary literature on Lewisian analysis, a number of commenters forward the claim that Lewis offers a two-level approach. For example, Livingston claims that Lewis:

“Envision[s] a program comprising two clearly distinct levels of analysis: first, the logical descriptions of causal roles, and second, the empirical identification of their occupiers.” (Livingston 2005, 35)

Likewise, Menzies and Price identify two distinct yet interconnected stages of Lewis’s program. They claim it is the second stage which satisfies Lewis’s metaphysical ambitions:

“The core of the second stage of Lewis’s program is that what the first stage provides, in effect, is a non-trivial target for empirical investigation: in this case, investigation of what it is, in fact, that plays the causal role R.” (Menzies & Price 2009, 6)

Whilst I think it is sensible to separate the analytical part of the project from the metaphysical part, I think it is misleading to call both these parts the ‘analysis’. To illustrate why, consider again Lewis’s metaphysics of mind. His argument for psychophysical identification consists of two stages (Lewis 1994b, 303). The analytical stage: where we define that some mental state M is the occupant of the M-role. The empirical stage: where we establish that some physical state P is the occupant of the M-role. From these two stages together we arrive at the conclusion that $M = P$. The analysis, therefore, is only part of the theory: the part that says that some mental state is equivalent to some functional role, and this is the case world-invariantly. The other part of the metaphysical theory is the empirical claim that there is some physical state which realises this functional role, which is specific to the actual world. The important part of Lewis’s program that gets us from analysis to ontological commitment is identifying what thing plays the functional role we have defined by analysis.

To strengthen my claims here, and relate analysis back to its primary objective, recall that I said that analysis need not rule out semantic indecision. Analysis

can still be useful even when it reveals that there is nothing that satisfies it perfectly. I take it that this commitment is in part motivated by the metaphysical work to which Lewis puts analysis. The process by which we define some theoretical role is quite traditional; however, the analysis has still done good work for us even if nothing perfectly satisfies the definition. Lewis's commitment to semantic indecision helps to ameliorate the distance between analysis and metaphysics. However, although this commitment enables analysis and metaphysics to work together, it should also help us appreciate why they cannot be the same. On the metaphysical side we are trying to provide some kind of generalisation, which does not have to be stateable or finite; by contrast, the analysis does have to be stateable and finite. Therefore, given that the world may imperfectly satisfy the analysis, any specification of that world cannot be given just by the analytic theory alone. As I have stressed: more must be done.

I have emphasized that Lewis thinks that more must be done to get you from a serviceable analysis to ontological commitments. There are two important steps. First: we must establish the non-emptiness of the analysis. That is: whether *anything* satisfies the analysis. Second: we must establish the character or nature of what satisfies the analysis. Of course, the analysis itself may tell us what the thing is like in causal terms; so this second stage may consist in a non-causal specification. In the application of the analytical results we must look further afield to provide the full metaphysical theory.

To summarise: I think we should interpret Lewis's conception of the analysis-metaphysics relationship in the following way. If we are going to try and give a metaphysical account of some phenomena, we require some idea of what would count as a satisfactory answer; I take it that analysis can be of great help here. It is hard to imagine exactly how we could even begin our metaphysical theorising had we not (to some extent) already established what the relevant terms of the discourse mean and how they are related. However, although the analysis clarifies the subject matter, this does not mean that the analysis is identical to the metaphysical account of that subject matter. The analytical process may indeed be a prerequisite for metaphysical theorising, but it is not the metaphysical theorising itself. In slogan form: Lewis takes it that analysis is

necessary, but insufficient for an adequate metaphysical methodology. However, although these projects are distinct in that they have different aims, in practice Lewis thinks that we should be concerned with both the analytical project and the metaphysical project given their methodological connection. (Lewis 1986b, 15)

v. Semantic Analysis and Metaphysical Reduction

In this section I want to discuss the claim that analysis provides us with a means of generating metaphysical reductions. I take this to forward the cause of the previous section: demonstrating how analysis and metaphysical reduction come apart further supports my claim that analysis and metaphysics are distinct theoretical enterprises.

I take this hypothesis to be an important one as some of the secondary literature use the terms ‘analysis’ and ‘reduction’ interchangeably¹⁰. In principle, I do not object to this, provided this terminological move does not obscure the sort of nuanced treatment of analysis I think is necessary. However, as the present proposal is concerned with *metaphysical reductions*, I think using ‘analysis’ and ‘metaphysical reduction’ interchangeably is misleading. Given that analysis cannot alone determine metaphysical commitments, analysis could not fully specify the reduction. That is: merely on the basis of the analytic theory, we will frequently not be in a position to identify how/whether the metaphysical base satisfies the conditions specified by the analysans. So analysis cannot be sufficient for metaphysical reduction as a reduction requires that there is a reductive base, and analysis alone cannot guarantee that.¹¹

Perhaps we could understand the proposal that analysis generates metaphysical reductions as suggesting that one aim for analysis is to reduce the number of primitives to which a theory must be committed. This does seem consistent

¹⁰ E.g. see Berkovski (2011, 102) and Maguire (2013, 133).

¹¹ For example, it doesn't seem implausible to think that a successful analysis may be consistent with an error theory.

with much of what Lewis says. Analysis does aim to reduce primitives. For example, Lewis tells us that: “The object of analysis is to reduce our burden of primitive notions, and to make tacit understanding explicit”. (Lewis 1986b, 154) This invites a question: what sort of primitives? Should the primitives that analysis aims to reduce be interpreted as primitive ontology, primitive ideology, or semantically primitive predicates? I argue that analysis *aims* to reduce semantic primitives.

You might be tempted to think that I am over complicating the story by introducing these distinctions. Why not think that in reducing the number of semantic primitives we thereby reduce metaphysical primitives? First, there is textual evidence to suggest that Lewis thinks that they come apart. Consider Lewis’s remarks in *On the Plurality of Worlds* regarding the balance of ontological and analytic questions:

“But if our work is directed to ontological questions and analytic questions both, as I think it usually is, then we are trying at once to cut down on questionable ontology and to cut down on primitives, and it is fair to object if one goal is served at too much a cost to the other.”
(Lewis 1986b, 15)

In opposing the acceptance of primitives against the acceptance of certain ontologies, Lewis gestures at the difference between *semantically primitive predicates* and being ontologically primitive or basic.

The most compelling reason to maintain this distinction, however, is that without it important theoretical subtlety will be ignored. Consider two examples. Presumably, if God exists, his existence is not dependent on anything. His existence does not (non-trivially) supervene on anything else, it is not grounded in anything else, it is not caused by anything else. Metaphysically speaking, God is primitive. Consider the predicate ‘is God’. Is this semantically primitive? No. ‘Is God’ can be analysed in terms of ‘is benevolent’, ‘is omniscient’, ‘is omnipresent’, and ‘is omnipotent’. So, we have found a case where the semantically primitive and the metaphysically primitive come apart: where there is something metaphysically basic, that can nonetheless be semantically analysed. This thought is significant as if we thought that because

there is more analytical work to be done there must be more metaphysical work, then we would end up drawing incorrect conclusions about the nature of God's existence.

Consider another example. According to Lewis's analysis, modal predications such as: 'it is possible that Danny DeVito is six feet tall' are to be analysed in terms of what is true of Danny DeVito's relevantly similar counterparts at other worlds. Semantically, this sentence cannot be further analysed. However, it seems that there is rather more we can say metaphysically. There will be cases where the semantically primitive is not metaphysically primitive. Whilst a semantic analysis of modal predications is going to bottom out at a counterpart relation in terms of similarity, and quantification over worlds, there is a much richer metaphysical story to be told pertaining to those worlds and counterpart relations. We may agree that the semantic analysis commits us to possible worlds, without having committed to a particular view of the nature of those worlds. I take Lewis's project in *On the Plurality of Worlds* to be engaged with providing that kind of metaphysical substance. Important metaphysical work would be ignored if we assumed that because the analysis of such sentences is complete that there is nothing metaphysically interesting left to say. This claim is well supported by previous discussion. Given that we know that analysis must be finitely stateable, whilst metaphysical commitments need not be, it seems reasonable to hold that substantial metaphysically reductive work will remain even once we have reached a semantic primitive.

Having said this, however, there is a sense in which analysis can *assist* metaphysical reduction. Analysing primitive modality in terms of quantification over worlds certainly reduces our primitive metaphysically ideological commitments, even if there is still ontological work left to do. There is also another way that analysis can play a role in metaphysical reduction. Consider the picture I laid out above regarding Lewis's metaphysics of mind. When we show that $M = P$, we have shown that M can be reduced to P. Nonetheless, the analysis of M as the occupant of the M-role provides a framework for the empirical investigation which identifies P as the M-role. In this way, analysis does provide us with a method for generating metaphysical reductions, albeit an incomplete one.

I think it best to characterise Lewis's view on analysis and metaphysical reduction as twofold. First there is the understanding of analysis as the shrinking of primitives. This should be understood as the shrinking of semantically primitive predicates. Sometimes, this will correspond to a shrinking of primitive ideology, and so can be said to aid metaphysical reduction. Second, analysis can be viewed as helping to provide ontological reductions. However, if only the analytical theory is known, such reductions are incomplete.

IV. Metaphysics and Truth

In the previous section I argued that analysis and metaphysics must be understood as distinct theoretical enterprises. The primary purpose of philosophical analysis is to give the meanings of the concepts under investigation. Whilst this enterprise is a prerequisite for successful metaphysical theorising, it is not identical to it. I claimed that the reason for this is that more needs to be done to arrive at the metaphysical theory. In this section I want to investigate the relationship between truth and metaphysics in Lewis's framework. To this end I will answer the following two questions. First, what role does truth play in metaphysics? Second, given this role, how should we understand the method and status of the resultant metaphysical enterprise?

vi. The Role of Truth in Metaphysics

I started this chapter by claiming that Lewis promises robust conceptions of both meaning and metaphysics, with a central and integrated role for truth in both these projects. To close this chapter I want to consider whether Lewis's project delivers on all it promised. What must be identified, therefore, are the salient and ineliminable places where true plays a role.

What role does Lewis assign to truth in the theory of meaning? There are two places where a role can be identified. First, in Lewis's project of radical interpretation, truth figures significantly in assigning an interpretation to the object language in the metalanguage. The output of this project gives the truth conditions for speakers' utterances. A second place where truth can be seen to have a role is in semantic analysis. Whilst both radical interpretation and semantic analysis aim to give meanings, the target of semantic analysis is different. Rather than aiming to give truth conditions for sentences, it aims to provide definitions for concepts. As previously described, an analysis is given by specifying the functional role of the concept under investigation, where this functional role is in part to be identified by systematising the output from the radical interpretation project. Lewis thinks that in giving a definition of the concept, analysis must provide world invariant truth conditions.

It is important to here to consider whether truth plays an ineliminable role in semantic analysis. If the aim is to specify the functional role of the concept under investigation, why does this need to be done in terms of truth conditions: couldn't the relevant roles be specified purely in the object language? There is of course room in theoretical space for a project which aims to specify the functional role of concepts without utilising truth conditions. It should be noted though that even if we take analysis to not be concerned with providing truth conditions, truth will yet have a role to play in analysis: not as an output of the approach, but as an input makes the approach possible.

Whilst this conception of semantic analysis, which is prosecuted purely in the object language, appears to be legitimate, we should ask whether this is Lewis's conception. That is: did Lewis think it necessary for analysis to provide truth conditions? I think it would be over committing to claim that Lewis thought that the analytic component of his approach *could not* be executed without utilising the notion of truth conditions. He doesn't really consider this as an option. Nonetheless, I think that the textual evidence supplied in section (i) provides a modest affirmative answer to this question of whether analysis should provide truth conditions – when giving the meaning, of either sentences or concepts, Lewis is explicit that giving the meaning is to be done in terms of truth conditions. (Lewis 1970b, 190) So not only will truth play a

facilitating role, with respect to the input for semantic analysis, but additionally it will feature as an output analysis.

Now to the question of the role assigned to truth in metaphysics. First, as we saw in (iii), Lewis is happy to endorse a general truthmaking claim: that truth supervenes on being. Second, if we take analysis to provide truth conditions, as I think we should, then, given that semantic analysis is a prerequisite for metaphysical inquiry, we have already identified a central role for truth in metaphysics. Part of metaphysical inquiry consists in finding the occupants of the roles identified by analysis, and if the specification of these roles is in terms of world invariant truth conditions, then truth is required to play an important role in the method of metaphysics.

Perhaps the identification of this role for truth in metaphysics can in turn provide some justification for my claim above that we should take truth conditions as the output from semantic analysis. Only some of the functional roles identified by analysis are going to place demands on reality – i.e. only the ones where the conditions specified by analysis are satisfied. There may well be many concepts which we can analyse for which there are no worldly satisfiers. But, when those conditions are satisfied, that tells us that there is something in the world which plays that role. As we saw in (iv), the non-emptiness of an analysis can be established by the relevant terms appearing in a true statement or theory. Thus, when a sentence (or theory) which contains the concept under investigation is true, this tells us that the x-conditions identified by analysis are satisfied. It seems reasonable, therefore, to call these conditions truth conditions.

The concept of truth that metaphysics inherits from the theory of meaning can only get us so far, inasmuch as the metaphysics is concerned. This limitation is well illustrated in (iv) where we see that, even after successful analysis, we can still ask questions about what the role players are: what is the character of the things to which analysis commits us. So how should one respond to this limitation? For all I've said about the role of truth in metaphysics, we might think that this is compatible with a much more minimal metaphysical enterprise. Perhaps we should cease metaphysical inquiry at the end of analysis, and accept that we cannot answer questions about the character of the role

players. We semantically analyse the true sentences of the language and then this gives us all the metaphysics we can get.¹² This is not Lewis's desired conception of the limit of metaphysics.

If we want a more robust conception of metaphysical inquiry, perhaps we could enrich the conception of truth with which we are working. Perhaps we need to place further restrictions on the specification of truth conditions such that when they are satisfied this settles all that is of metaphysical interest.¹³ I suggest that given the constraints that Lewis places on the concept of truth he cannot do this. The truth predicate for Lewis is a theoretical term of the semantic theory. If Lewis tried to offer 'metaphysical truth conditions' it isn't that these would be falsehoods, but rather such 'truth conditions' would be gratuitous. In order to express what we want to in fundamental metaphysics, we may as well say: p iff s . The introduction of 'is true' at this level isn't doing any explanatory work for Lewis, and as such he should not pursue this move. I've already argued Lewis cannot adopt a truthmaker conception of truth as it would be in tension with his aims for and constraints on analysis.

So two questions for Lewis: how far does truth get us, and how does he respond to the limit? I argue that truth gets us pretty far, as the project of semantic analysis gives us quite a lot of direction with respect to metaphysical commitments. Metaphysics is thus well constrained on the Lewisian picture. However, there is still much of metaphysical interest left unsettled. By way of example, recall Lewis's remarks about how he can agree with the Ersatzist about the analysis of modality but radically diverge with respect to the metaphysics. Thus, if we want a robust conception of metaphysics which doesn't place any further demands on the conception of truth we need to employ independent metaphysical methods to determine the nature of the role players. So Lewis does deliver on a robust conception of metaphysics, but this project, whilst constrained, is autonomous. My central claim about Lewis and the role of truth in metaphysics is that the truth does not have work to do

¹² I consider options like this in Chapters Four and Five.

¹³ I consider an option like this in Chapter Three.

beyond its role in semantic analysis and the broad claim that truth supervenes on being.

vii. The Metaphysical Enterprise

In this section I want to elucidate the significant remaining questions for Lewis. I have argued that if we go with Lewis's conception of the role of truth in metaphysics, then we need to employ independent metaphysical methods to get us beyond the limit of semantic analysis. This invites the following questions: (1) What are these methods? (2) Are these methods legitimate? (3) Are these methods truth conducive?

The answer to (1) is twofold. First, given the broadly naturalistic outlook of Lewis's metaphysics, in many cases the methods used to determine the nature of the role players just are the methods of natural science. Think again about the psychophysical identification. Identifying the worldly role players of the M-Role is work for neurobiologists, and thus requires the best methods of that scientific discipline. Second, where natural science will not deliver a verdict on the nature of the role players Lewis adopts a cost-benefit analysis involving theoretical virtues.¹⁴ This typically involves assessing a candidate theory for parsimony, explanatory power, elegance, applicability, coherence with other commitments, etc. The virtues (and vices) of this theory are then weighed in contrast to other theories of the same subject matter. The theory with the greatest virtues on balance is the theory which should be adopted.

This naturally invites questions (2) and (3). Why should such methods be accepted, and why should we think that these methods will yield true beliefs about the world? The questions are obviously related, showing that (3) is the case would provide justification for (2). That is: if it can be shown that such virtues do yield true belief about the world, then this would provide a motivation for a method that utilises them. However, even if (2) can be motivated because theories which possess such virtues make theories more

¹⁴I take Lewis (1986b) to provide the most comprehensive description of this method.

intelligible or computable for creatures like us, we might still wonder whether the possession of these features is truth conducive.

Definitively answering (2) and (3) falls outside of the scope of my thesis. Nonetheless, merely articulating these questions and the relationship between them is useful as it illuminates what would be required to adopt Lewis's approach to metaphysics. If an affirmative answer to (2) could be achieved, but not to (3), then this suggests a conception of metaphysics which may well depart from Lewis's robustly realist ambitions. However, if an affirmative answer to (3), and consequently (2), could be achieved then the adoption of this method for metaphysics promises a realist enterprise. It is the latter of these options that seems more in keeping with how Lewis describes the aims and statuses of the metaphysical theories he offers. Of course the argument needs to be made. Therefore, if you favour an approach to metaphysics which proceeds via semantic analysis in the way Lewis endorses and want that approach to deliver a robust, realist metaphysics, then what Lewis offers is going to be conditional upon providing satisfactory answers to these questions.

V. Conclusion

In this chapter I offered a critical evaluation of David Lewis's project of philosophical analysis. I have argued that analysis and metaphysics must be understood as distinct theoretical enterprises. The primary purpose of philosophical analysis is to give the meanings of the concepts under investigation. Whilst this enterprise is a prerequisite for successful metaphysical theorising, it is not identical to it. My central claim about Lewis and the role of truth in metaphysics is that the concept of truth does not have work to do beyond its role in semantic analysis and the broad supervenience claim. Adopting Lewis's approach therefore requires utilising distinct metaphysical methods to answer metaphysical questions, and thus the viability of the project rests on the legitimacy of these methods.

Chapter Two: Heather Dyke and Unmediated Metaphysics

I. Introduction

In *Metaphysics and the Representational Fallacy* (2008) Heather Dyke presents a methodological challenge to a certain way of doing metaphysics. The aim of her book is to argue against what she terms ‘the representational fallacy’, which is “a general strategy of reading metaphysics off language.” (Dyke 2008, 7) In place of this allegedly problematic strategy, she proposes an alternative theoretical framework in which to prosecute serious metaphysical inquiry.

This chapter offers a critique of Dyke’s proposal. The proposal is representative of a certain animosity towards taking meaning and language to be in some sense authoritative when doing metaphysics.¹ My present investigation therefore takes Dyke’s position as a case study for the prospects of an unmediated approach to metaphysics. I argue both that Dyke’s arguments against taking language seriously when doing metaphysics, and her proposed alternative methodological strategy are unsuccessful. With respect to the former, I contend that her arguments do not have the far reaching methodological significance she intends; with respect to the latter, I argue that her proposal is either subject to some of her own criticisms, or that it presupposes the use of theoretical tools which are left both underspecified and undermotivated.

The chapter proceeds as follows. In §II, I present Dyke’s proposal. First, I situate her project and sketch her view of the aims of metaphysics. Second, I outline the methodology that Dyke takes to be problematic and her argument against it. Third, I present her preferred methodological strategy. In §III, I turn

¹ Similar sentiments can be found in John Heil (2003) and Michael Devitt’s (1983) work.

to offer criticisms of Dyke's proposal. First, I argue that Dyke's preferred method problematically commits her to the view that all metaphysics is realist metaphysics. Second, I demonstrate that her own proposal falls short of her methodological aspirations: either her arguments are not fatal to the metaphysician who wants to take language seriously, or they end up being self-defeating. Finally, and most importantly, I consider the role she assigns to truth in her theoretical framework, and argue that she cannot sustain this role without undermining her proposed theoretical framework. In §IV, I conclude by offering what I take to be the generalizable lessons of the preceding discussion.

II. The Representational Fallacy – Overview

i. Situating the Project

Dyke offers a historical reconstruction of the development of metaphysics, and presents this as a partial explanation for why metaphysicians are systematically methodologically misguided. Her story starts from Kant, through to the logical positivists. Dyke argues that these philosophers did not themselves commit the representational fallacy, as they did not think metaphysics (what we might call realist, or perhaps noumenal metaphysics) was a possible subject of philosophical investigation. These philosophers thought that there are insurmountable epistemic or linguistic barriers to theorising about the world as it is in and of itself. Philosophy, they argued, must therefore content itself with a different kind of theoretical enterprise.

Dyke then argues that the 'linguistic turn' of the positivists was followed by a re-emergence of metaphysics, but that this re-emergence did not satisfactorily shake off the prior preoccupation with language. In using the tools of a philosophy which has decidedly non-metaphysical objects and aims, these post-turn metaphysicians are, the argument goes, ill-equipped to deal with the

demands of a thoroughgoing realist metaphysics. These philosophers maintained that the proper focus of philosophical inquiry is language, but they *also* thought that extralinguistic reality avails itself to substantive investigation. Consequently, they believed that “the study of language must be the proper route to knowledge of it”. (Dyke 2008, 21) Dyke makes the bold claim that the project of 20th and 21st Century metaphysics has been to try and strive for the ‘ideal’ language in which to describe extralinguistic reality. (Dyke 2008, 8-9)

Setting the question of historical accuracy aside, it should be clear that there are some ways of viewing the aim of metaphysical inquiry such that starting with language is entirely appropriate. For example, if I think that the aim of metaphysics is to understand and systematise the conceptual repertoire needed to think about and navigate extralinguistic reality, then starting with our linguistic representations of reality seems apt. However, for Dyke, this is not the proper aim of metaphysics, and a tension arises when we think *both* that language is the proper source of philosophical knowledge and when we also think that the aim of inquiry is a theory about reality as it is in and of itself. Dyke holds that only the latter of these views is defensible. She tells us that:

“Metaphysics studies the nature of existence, what kind (or kinds) of things exist. Its subject matter is reality itself, and its aim is to discover the nature and structure of that reality.”² (Dyke 2008, 22)

She positions Aristotelianism as the desired metaphysical project. (Dyke 2008, 22-23) She claims that the purpose of metaphysical inquiry is to uncover the fundamental structure of reality. She is careful to distinguish between a metaphysically deflationary conception of this project, according to which metaphysics is exhausted by the inquiry of natural science, and the view she prefers, which she considers to be a robust realist conception of metaphysics. She claims that we can be robust realist metaphysicians whilst consistently “denying that we can learn about the nature or structure of reality from the

² You might be concerned that this way of putting the point glides over some substantial issues about whether all being reduces to existence. It seems like a serious metaphysical question whether ideology reduces to ontology. To say that metaphysics studies existence seems to worryingly neglect metaphysical views which think that metaphysics studies what exists and what those things that exist are like. Such views take the latter question to be irreducibly non-existential. Cf. Melia (1995).

nature or structure of the language we use to describe it.” (Dyke 2008, 21) Her project is thus an antidote to what she takes to be the misguided metaphysics of the 20th and 21st Century.

ii. The Problematic Methodology

Here is a crude sketch of how we might think we should approach metaphysical inquiry. We start with some kind of general existential question, we then examine our pre-theoretic thought and talk about that subject matter and use this to determine the correct answer to the existential question. (Dyke 2008, 27-28) If we do talk about the subject matter (i.e. seem to utter true sentences about it) but we don’t want to affirm the existential question, we can perhaps offer a paraphrase, where we try to do away with the seeming commitment. Dyke describes this as “A typical methodology... [which] involves taking our prephilosophical way of representing reality and asking what sorts of entities their truth implies the existence of.” (Dyke 2008, 9)

Dyke claims that philosophers are frequently seduced into theorising as above, because they implicitly endorse a strong view about the relationship between language and the world. The assumption that gets us from a linguistic claim about a certain portion of the language (i.e. that the sentences that make up that portion of the language are ineliminable and true), to a view about corresponding ontological facts about what the world is like, is that there is one true description of reality. (Dyke 2008, 45) She calls this the ‘strong linguistic thesis’:

“SLT: There is one privileged, true description of reality, the sentences of which (a) stand in a one-to-one correspondence with facts in the world, and (b) are structurally isomorphic to the facts with which they correspond.” (Dyke 2008, 46)

The upshot of a commitment to SLT is that putatively ontological debates get caught up in arguing about the meanings of the suspect terms. If we have true sentences which contain the suspect terms, either those sentences themselves

will be part of the one true description, or they must be paraphrased away, in favour of synonymous sentences which do not contain the suspect terms.

There is a dual concern here. First, metaphysicians will get caught up merely arguing about semantic issues, and so may not get to the serious business of metaphysics at all. But, second, even if they do, this initial preoccupation with meaning and translation will skew any subsequent debate such that the resultant theories will not be able to satisfy their Aristotelian ambitions.

Dyke is explicit that she thinks that SLT is false. (Dyke 2008, 71) Her master argument against this approach is that many different ontologies are compatible with a particular language. As she says: “All that is central to my view is that two or more truths can have the same truthmaker while being non-synonymous.” (Dyke 2008, 99) The view underpinning this is that there are multiple accurate ways of describing reality, which may not be reducible to each other. (Dyke 2008, 99) Dyke wants to deny the claim, therefore, that all true sentences will be reducible to one privileged class of truths. (Dyke 2008, 72) This is connected to the method of paraphrase. She doesn’t like the thought that the terms which appear in the privileged description only do so because we are unable to paraphrase them away. (Dyke 2008, 72) By way of example for how the SLT functions, she offers indispensability arguments for mathematical Platonism: she claims that those who endorse indispensability arguments must assume that the fact that mathematical discourse is ineliminable from true scientific discourse entails something about the truthmakers for such discourse. (Dyke 2008, 72)

It must be noted that Dyke takes her argument against SLT to be compatible with the claim that some descriptions can be better than others; that they can say ‘more objectively’ what reality is like. (Dyke 2008, 71) For example, she claims that in theories of fundamental physics we do seem to have correspondence between the description of reality, and reality itself. (Dyke 2008, 75) Dyke maintains, however, that this claim does not commit one to the thesis that there is *one* true description of reality. Even if there is a metaphysically perspicuous way to describe the facts of fundamental physics, it does not mean that those facts won’t also serve as truthmakers for metaphysically unperspicuous true sentences. She thus holds that metaphysics

should not be in the business of reading any metaphysical commitments off the true sentences of our language.

iii. The Proposed Alternative

Dyke claims that theorising in line with the SLT results in an inadequate conception of the theoretical options. Once it has been rejected, however, a new position emerges. She calls this the ‘overlooked strategy’. Broadly, this strategy is to shift metaphysical debate from a question of whether ‘there are Fs’ is true, to a question of whether Fs are among the truthmakers for true sentences. She thinks once we make this move we will then be liberated to consider the advantages and disadvantages of the metaphysical theories themselves. (Dyke 2008, 88) She outlines this new strategy as follows. First, we “explore the possible ways the world might be, ruling out of contention any ways that it is not possible for the world to be”, and then we “argue that one of these way the world might be is the way the world actually is.” (Dyke 2008, 60) The former is a distinctly metaphysical stage, and normally employs a priori reasoning. The second stage needs to appeal to empirical evidence, and so “in this respect metaphysics is continuous with science.” (Dyke 2008, 60)

I want to briefly gesture at an expositional concern. Dyke doesn’t give much by way of example for how this methodology is supposed to unfold. She tells us in the case of the metaphysics of time that stage one consists in demonstrating the incoherence of A-theory, which shows why it cannot be a genuine metaphysical possibility. The second stage consists in arguing that B-theory is the best theory with respect to our best confirmed current scientific theory: that B-theory is consistent with the special theory of relativity. I think given the straight-to-metaphysics approach she is advocating, there is a serious question about what kind of incoherence is being demonstrated in this first stage. You might think that it must be conceptual incoherence, but then it seems like she will be engaging in taking conceptual relations to tell us something about the world; a move I take it she wouldn’t want. If it isn’t conceptual incoherence then perhaps it is that A-theory (on some interpretations) is inconsistent with

the theory of special relativity. However, this suggests that stage one and two collapse into each other. Working out the details of the proposal is important. However, given that the primary critical concern of this chapter is to discern whether Dyke's method assigns a legitimate role to truth, I will proceed by assuming that these details can be worked out in a satisfactory way.

When discussing how this overlooked strategy relates to the aims of metaphysics Dyke claims that "the aim of metaphysics is to provide us with a map of the structure of all that exists." (Dyke 2008, 63) Whilst this way of phrasing the point may make it a little opaque, I take it that she means that in order to construct a map we need to start from the terrain. Dyke explicitly criticises 'top-down' methodology: one that starts with true representations and infers ontology from those representations. (Dyke 2008, 73) In addition, another thought that seems consistent with what she says in these passages is that we may have different maps for different explanatory purposes. So, starting with the maps may tell us more about our explanatory needs, than about what the world is like independently of them. To put her argument against SLT in these terms: multiple maps are going to be consistent with a single terrain.

Given that truthmakers are going to be playing a central role in Dyke's metaphysics, we need a better idea of how it is that she is thinking about them. Dyke defines *facts* as follows:

"I take a fact to be simply a portion of reality, a bit of the world, a way some part of reality is. I am, thus, not using the term to denote a linguistic entity: a true sentence or proposition. I am, instead, using it to denote an extralinguistic entity: a part of the world." (Dyke 2008, 39)

She claims that she means to use this understanding of 'facts' and 'truthmaker' interchangeably. (Dyke 2008, 39) Dyke, however, does not offer a fully developed theory of truthmaking or truthmakers. (Dyke 2008, 77) She does not think that the method she offers requires any one particular view about what truthmakers are, but rather has a kind of generality. (Dyke 2008, 81) She claims that we have an intuitive understanding of what locutions like 'making true' or 'true in virtue of' mean, and she is happy to stick with this intuitive

appreciation. (Dyke 2008, 79) This approach has merit. If the aim of her proposal is to provide a theoretical framework in which to engage metaphysical questions, then making the proposal neutral with respect to different theories of truthmakers may make her suggestion more appealing.

Dyke claims that once we appreciate the overlooked strategy, and we have our intuitive grasp of truthmakers and truthmaking, we see that “the truthmaking project undermines the paraphrasing project.” (Dyke 2008, 85) In the paraphrase project we sought to find acceptable translations of sentences containing suspect terms into sentences not containing such terms. The thought being that our true unparaphrasable sentences are going to determine our metaphysical commitments. However, once we move to a conception of metaphysical theorising that is primarily concerned with truthmaking, we see that multiple non-synonymous sentences can have the same truthmakers, and so we should not paraphrase anything away. As Dyke says: “The alternative position that has been overlooked is that truths from outside the realm of physics are made true by physical facts, facts describable in the language of physics, even though they are not reducible to physical truths.”³ (Dyke 2008, 98)

III. The Representational Fallacy – Critique

In this section I offer a critique of Dyke’s project. The critique consists of three stages. First, I evaluate her view of the status of metaphysics. I argue that Dyke’s suggested aim for metaphysics, and her preferred method for metaphysics, problematically commit her to the view that all metaphysics is realist metaphysics. Second, I evaluate her argument against taking language to in some sense be authoritative when doing metaphysics. I argue that either her arguments are not fatal to the metaphysician who wants to take language

³ It’s not clear to me that this has actually been overlooked. This seems like the method pursued by Lewis (1994) with respect to psychophysical identification.

seriously, or they end up being self-defeating. Third, I evaluate the role assigned to truth in her framework. I show how her proposal requires that there is a conception of truth available which is absent of the demands placed on the concept by the theory of meaning. I then argue that she cannot sustain this role without undermining her proposed theoretical framework.

iv. All Metaphysics as Realist Metaphysics

Dyke claims that the purpose of metaphysical inquiry is to uncover the fundamental structure of reality. You might think that there is a substantial and significant distinction between a conception of metaphysical inquiry according to which realism is an option, and one according to which realist conclusions are the only possible ones. However, given the way that Dyke sets up the project of metaphysics and the resources that she allows herself, she cannot make sense of such a distinction and her project consequently precludes tout court many well established anti-realist positions. The challenge I forward in this section is essentially a methodological one: because of how Dyke sets up the aims and methods of metaphysics, she deprives metametaphysics of some important theoretical options.

Let's start by considering Dyke's historical reconstruction of the development of metaphysics. Once we've moved away from Kantianism, (or indeed perhaps other kinds of metametaphysical views such as Berkeley's idealism or Strawson's descriptive metaphysics) we recover metaphysical realism as an option. It is therefore an open question whether we endorse realism or anti-realism about a particular class of entities, properties, or relations. So, instead of us having to buy wholesale one view or the other, we can proceed in a piecemeal way to work out the ontological status of a particular class of entities, properties, or relations. For example, we may adopt a robust realism about the posits of physics, projectivism about moral properties, and a response dependent theory of colour. In each case, there is serious metaphysics to be done, but the resultant theories make very different claims about the nature of the entities or properties in question.

However, if the only proper aim of metaphysics is to provide a robustly realist theory of the fundamental structure of reality, then all theories which are not in the business of doing that will not count as proper metaphysics. It seems at least *prima facie* odd that Dyke cannot accommodate a practicable distinction between having realism as an option, versus having realism as the only option. If we can do metaphysics at all it seems we should, in principle, be able to have a genuine independent debate about the ontological status of any putative entities, properties, or relations to which we are committed. However, merely from her characterisation of the aims of metaphysics it is not clear how such debate can be executed. Dyke sets up the aims of metaphysics as full-blooded and robust, and ties both of these commitments to a conception of metaphysics as realist.⁴

In addition to the way she articulates the aims of metaphysics, Dyke's method itself suggests that the distinction cannot be maintained. Recall her positive proposal. When engaging in metaphysical inquiry, we start by considering the range of ways the world might be. Of those possible ways, we then discover what the truthmakers are, and this settles the way the world actually is. Knowing any kind of semantic information won't help us here, as according to Dyke, knowing that a sentence is true, "even knowing what its truth condition is, does not tell us what it is about the world that makes it true." (Dyke 2008, 65) Metaphysics sole task is to discover what the truthmakers are, and in virtue of these extra-linguistic facts the sentences of the language are made true. Naturally, there are going to be false sentences. There will be objects and properties we appear to talk about for which there are no truthmakers. So, Dyke's picture may be able to accommodate error-theoretic anti-realism. Prior to metaphysical inquiry we may think certain kinds of sentences are true, but once we have discovered what the truthmakers are we realise that actually a class of sentences are false.

⁴ You might also worry about how such a characterisation of the aims of metaphysics would exclude metaphysical inquiry into so-called 'derivative ontology'. If the only proper aim of metaphysics is to provide theories of fundamental reality, then inquiry into the existence and nature of, for example, social kinds would also be off the table. For some interesting discussion of the metametaphysical significance of the nonfundamental, see Barnes (2014).

What of other forms of anti-realism? Dyke tells us that non-cognitivism is “unappealing”. (Dyke 2008, 130) However, at the very least, even if you don’t want to endorse a non-cognitivist view, the question of how sentences are made true or if they are even eligible for truth seems like a substantive philosophical question. There is more, it seems, to the relationship between language and metaphysics than just truthmaking. Even if we decide to agree with Dyke that truthmaking is the primary concern for metaphysics, being able to account for discourses which don’t aim to describe the fundamental nature of reality and give a full blooded and robust account of how these function should also be a job for metaphysics. Of course, Dyke may dig her heels in and insist that such a project is not one happily situated in metaphysics. She might claim that these sentiments are merely a throwback from an outdated way of thinking. However, I think this should give us pause. To claim this would be to rule out most of the classic and contemporary debate in the realist-anti-realist literature as not really metaphysics. What is worse, this exclusion is just a consequence of the way the method is set up, and not a product of any serious argument against the particular anti-realist positions. There seems something dialectically inappropriate in such a move, and I think Dyke owes these positions more.

Finally, there are ways of engaging with the question of anti-realism which do not solely attend to considerations to do with language and truth. One concerning consequence of Dyke’s view seems to be that even some prominent forms of Aristotelian metaphysics seem to be off the table. Kit Fine (2001) argues that we can use a neo-Aristotelian framework to determine whether we should be realist or anti-realist about some candidate entities, by attending to the question of what grounds what. Whilst Dyke is well within her rights to disagree with the Finean metaphysical program, it seems that on her view much of Fine’s approach would not count as metaphysics if anti-realism is ruled out. There is a way of making sense of the Finean view such that it is not committing the representational fallacy: we really are concerned with the fundamental facts, but something interesting nonetheless hangs on the relationship these facts bear to our ways of thinking about the world. However, if the only proper aim for metaphysics is something robustly realist, it is hard

to see how a Finean project which in part aims to explain how the fundamental relates to the nonfundamental can be accommodated.

I conclude, therefore, that Dyke's position carves up the theoretical space in a way such that there is an unacceptable loss of theoretical options, which is damaging to the plausibility of both her characterisation of the aims of metaphysics, and the methods she employs to satisfy those aims. This is all the more worrying given that the aim of her project is to provide a framework from which we can prosecute metaphysical theorising.

v. Truthmaking and the Representational Fallacy

In this section I argue that Dyke's proposed methodology itself falls short of her own theoretical aspirations. My critique proceeds as follows. First, I show that there are two different readings of Dyke's claim that we ought not take the relationship between language and world to be significant when doing metaphysics: a weak reading and a strong reading. I then argue that both readings severely compromise Dyke's position. On the weak reading, the claim doesn't have the far reaching methodological significance her rhetoric suggests. But on the strong reading, Dyke is committed either to a mysterious truthmaking relation, or to a relation which appears to be in conflict with the strong reading itself. Finally, I consider a possible concessionary move on Dyke's part: that she allow for a metaphysically significant relation between language and world, but one that must obey a certain direction of fit. I argue that this move cannot do the work that Dyke needs it to, as it fails to respect a crucial distinction between epistemic and metaphysical concerns.

Dyke offers us a 'metaphysics first' conception of metaphysical method. She maintains that something has gone wrong when we focus on the relationship between language and the world when doing metaphysics. Not only is metaphysics *about* the world, but the methods we use to theorise about it should not make central concerns to do with the relation between language and the world. Her charge is that those philosophers who take concerns to do

with truth as it functions in the theory of meaning, and language to be significant when doing metaphysics have made a mistake.

Now, there is a strong and a weak reading of some of Dyke's claims with regards to these matters. On the weak reading, philosophers have gone wrong because they have focussed their attention on the wrong relation. Perhaps they have been concerned with a method which only looks to the true sentences of the language to determine ontological commitments, rather than doing any independent metaphysical work. Perhaps they were misguided because they thought that there was only one such relation, when in fact there are many. This latter view in particular seems to be at the root of Dyke's master argument against SLT. There are many equally important relations between the language we use to describe the world and the world itself, and so metaphysics should not just focus on one such relation. But, there is a much stronger reading as well. According to this second reading philosophers have made a mistake in thinking that focussing on *any* such relation is appropriate work for metaphysicians. To put the point another way: according to the weak view *some* ways of methodologically proceeding from language to the world are wrong. According to the strong view, *any* method which moves from language to the world is mistaken.

Dyke seems to endorse latter of these readings but her arguments only really seems to generate, if anything at all, a reason to accept the former. I will consider the strong reading below, but first some comment on the weak reading. There is a gap between the rhetoric and the argument that Dyke forwards. Her master argument, if successful, merely shows that there is not a unique fit between language and world: that there are multiple ways that our language can hook up with reality. The trouble for Dyke is that it doesn't follow from this underdetermination that there could not be a privileged true description. She concedes that the language of physics is in some sense more accurate than other discourses. But if this is so why not think that the complete language of complete physics provides something like a privileged description of reality? If Dyke's proposed alternative method for metaphysics is to be practicable, I assume we should be able to name the truthmakers for the true sentences of our language. If we can't do this, then it is not clear how we could

begin to theorise about the truthmakers in the first place. Given this, couldn't the privileged language just be that language that names the truthmakers? Maybe there will be other languages which state truths about the world in a metaphysically unobtrusive ways, but there will be at least one language which is in the business of describing the fundamental structure of reality.

At the very least, even if we cannot recover a privileged description, if all Dyke aims to defeat is the view that there is a single unique relation between language and the world, then her attack doesn't seem so damning. Indeed, this would be consistent with numerous arguments for semantic indeterminacy. This is even compatible with the view, held by some metaphysicians, that when we are engaged in metaphysical inquiry we theorise using a special metaphysical language.⁵ But I take it that she intends more than this. She claims that her intention is to rid us of a way of doing metaphysics according to which word-world relations are relevant at all. It isn't merely that some conception of the relationship between language and world is off the table, but rather that we shouldn't be taking representation seriously at all. This is the strong claim. And if this is what she intends, then her own preferred method cannot satisfy it.

Dyke's alternative strategy consists in determine what the truthmakers are. Thus at the centre of the proposed method is a notion of truthmaking. As I noted, Dyke doesn't tell us much about what view of truthmaking and truthmakers we should prefer, but minimally it seems uncontroversial to claim that truthmaking is a relation. This invites a question. What are the relata? Now, perhaps it is a relation between facts (i.e. truthmakers) and something else non-representational. I have trouble making sense of this view. What could the non-representational relatum be? If Dyke's preferred method entails us committing to such a mysterious relation, then I think it is not unreasonable to claim that she has not given us a viable philosophical method. After all, it is far from clear how to proceed with a method where the central machinery of this method is in some sense mysterious.

Alternatively, perhaps it is a relation between something with representational content (a truth-bearer) and a fact. However, if this is the way to understand

⁵ For example, see Sider (2009 & 2011), Cameron (2010), and Dorr (2008).

truthmaking, and I think it is, then it seems there is something to the thought that some claims about the relationship between language and world are permissible in metaphysics after all. This is a long way from a commitment to SLT, but it is nevertheless a claim that there is a space in metaphysics for theorising about the relationship between our representations of the world and the world itself. Inasmuch as this is the case, this falls foul of the strong reading.

At this point Dyke might claim that that the concern about relations between representations and world is a concern about the direction of fit between the two. So, we can move from claims about the world, to claims about language safely, as this is the correct direction of explanation. What is impermissible is moving from claims about the true sentences of the language to claims about what the world is like. I am in agreement with Dyke that, “the direction of fit between language and reality is not from language to reality, but from reality to language.” (Dyke 2008, 66) But this can be read very weakly as no more than a commitment to a non-idealist metaphysics, where this consists in commitments to the following two claims: (a) there is a world independent of my representations of it, and (b) the majority of the true sentences I utter are true (to use the popular locution) *in virtue of* what the world is like. But a commitment to this (admittedly coarse-grained and radically underspecified) view of the relationship between language and metaphysics, tells us very little about the method I should employ in order to do metaphysics. I can accept this view and consistently think that methodologically I must proceed from systems of representation to facts about the world.

Perhaps Dyke would say that this gets the order of explanation wrong. Her constant claim is that we cannot read what the truthmakers are off sentences. (Dyke 2008, 57) Once we appreciate this, I take it she thinks that we can essentially reverse the direction of explanation between the semantics and the metaphysics. So, rather than moving from meaning facts to world facts (the representational fallacy), we first attend to the metaphysical facts which means that we can then explain why certain sentences of the language are true. (Dyke 2008, 54) In response to this I think the crucial observation is that to claim there is an order of metaphysical priority, is to say nothing of the order to

methodological priority. I take this as a familiar distinction in the philosophy of science, which is helpfully elucidated by Reichenbach:

“The *objective* relation from the given entities to the solution, and the *subjective* way of finding it, are clearly separated for problems of a deductive character; we must learn to make the same distinction for the problem of the inductive relation from facts to theories.” (Reichenbach 1938, 36– 37)

Dyke’s framework problematically blurs this metaphysical-epistemological distinction.

A defender of Dyke might claim that the fact that there isn’t a uniquely best description of reality shows that our epistemological route to knowledge of reality should not be in the business of utilising our representations of reality. If we take Dyke’s target to be the metaphysician who takes the true sentences of the language to be the only authoritative source of evidence in metaphysics, then this defence seems convincing. However, we can imagine a more modest metaphysician who takes such evidence to provide a defeasible source of evidence for theory construction in metaphysics. With this target in mind, the defence still seems successful in the case where we have a single sentence which could be made true by different incompatible truthmakers.

However, in the case where we have multiple non-synonymous sentences, all of which share a truthmaker, it just does not follow from the fact that a set of sentences all share a truthmaker that the study of those sentences cannot facilitate our investigation into what their truthmaker is. Of course, we will often not be in a position to know whether we are dealing with the former or the latter case, so the putative evidence should be approached with caution. But, given that Dyke aims to establish that we should never take the true sentences of the language to provide evidence for a metaphysical theory, I trust that the preceding discussion can provide some restitution for those who want a method for metaphysics which takes our representations of reality to provide some evidence for what reality is like.

vi. Motivating the Role of Truth

In this section I assess the legitimacy of Dyke's deployment of the notion of truth in her proposed method. The task at hand is to get a handle on how Dyke motivates truth's role in her theoretical framework. To clarify, I take it that Dyke's project is not in any way deflationary about truth: truth plays an indispensable role in both the theory of meaning and in metaphysics. The section consists of two stages. First, I evaluate what Dyke says about the relationship between the role assigned to truth in the theory of meaning and the role assigned to it in metaphysics. I aim to ascertain whether it is possible to motivate the role of truth in metaphysics via the role that it plays in the theory of meaning. Second, given that Dyke requires a concept of truth which is free from the constraints placed on it in the theory of meaning, I aim to determine whether this role for truth in metaphysics is can be motivated, and if so how.

Although "talk of truth conditions is commonplace throughout philosophy", Dyke claims that "the notion of truth conditions is ambiguous between semantics and ontology." (Dyke 2008, 48) We might think that talk of truth conditions happily sits in a project of giving the meaning for sentences of the language. Or, we might think that talk of truth conditions speaks to ontological concerns in that there is "a significant connection between the truth condition of a sentence and the ontology that grounds the truth of that sentence if it is true". (Dyke 2008, 49) It is important to note that even those philosophers who think that metaphysical inquiry is posterior to the satisfactory semantic analysis of language do not think that semantics alone can give us metaphysical commitments. Even if you think that language is in some sense the proper source of knowledge in metaphysics, you still need to know the truth values for the statements of a language, as well as the truth conditions. Of course Dyke thinks that even this isn't enough. We need to know the truthmakers, as well as truth values: merely knowing the truth conditions for a true sentence is not sufficient for telling us what in the world *makes* said sentence true.

Dyke holds that truth conditions ambiguously serve a dual function. There is an ontological function of providing truth conditions, which pertains to the relationship between truth and reality, and a semantic function, which pertains to the relationship between truth and meaning. (Dyke 2008, 50) Dyke thinks that to acknowledge that there is this dual theoretical role for truth conditions is not to commit the representational fallacy; it all depends on the work to which the truth conditions are being put. (Dyke 2008, 53) I think there are some serious exegetical concerns here about exactly how this ambiguity view works. Is the thought that irrespective of whether truth conditions are put to semantic or metaphysical work, what ‘truth conditions’ means is stable across both of deployments? But then, depending on whether they are put to ontological or semantic work, the manner in which the truth conditions are expressed can be subject to change? So when talk of truth conditions is put to ontological work we have some story about what constrains their use in that context, and when they are put to their semantic work we also have a story about what constrains their use on this different context. The form this might take is that there is some kind of restriction placed on the specification of truth conditions when they are put to metaphysical work. For example, perhaps the right-hand side of the biconditional must be stated in joint carving terms.⁶ And when they are put to semantic work the specifications of the truth conditions must be constrained by what speakers need to know if they are to understand a sentence? Then we commit the representational fallacy when we infer something about the nature of reality from truth conditions which are serving their semantic function.

If Dyke really does think that talk of truth conditions is ambiguous, then it would be good to see how these supposed two different functions relate, if they are supposed to at all. For example, imagine we give the (semantic) truth conditions for a sentence of the language and in so doing give the meaning of that sentence. Let’s say that that sentence is true. We then want to hold that there is a significant connection between the (metaphysical) truth condition of that sentence and the ontology that grounds the truth of that sentence. So what

⁶ I consider this move in the next chapter.

are the truthmakers making true? Are the truthmakers satisfying (1) the (semantic) truth conditions, (2) the (metaphysical) truth conditions, or (3) both? If (1) or (3) then it is hard to see what theoretical work the (metaphysical) truth conditions do. It seems like we need a story here about what theoretical role the metaphysically functioning truth conditions play, that isn't already captured by the talk of truthmaking. If (2) then it seems very unclear under what conditions we could appropriately talk about (semantic) truth conditions being satisfied. Further, I presume that we cannot translate the (semantic) truth condition into the (metaphysical) truth condition, as this would be to prosecute a paraphrase project. The root of this concern is that once we try to flesh out the story about how metaphysics relates to the theory of meaning we might find that whatever sense in which the metaphysical theory is 'true' does not appropriately support the notion at work in the theory of meaning. However, as it is not Dyke's intention to provide a theory of meaning, it would be a misplaced criticism to claim that her framework does not satisfactorily explain the relationship between the theory of meaning and truth in metaphysics.⁷

What I hope to have established thus far is that whatever is motivating Dyke's talk of truth and correspondence in her project, this motivation does not come from any reflections on the role of truth in the theory of meaning. So, with respect to metaphysics, what does Dyke tell us about truth? First, she tells us that our true sentences are true because of the way the world is. This seems to make some use of a notion of correspondence. Our systems of representation don't float free from the world: they can 'get it right'. Second, and perhaps most importantly, her central theoretical tools are truthmakers. Both the talk of truth as correspondence and truthmakers invites us to consider what legitimises Dyke's deployment of these notions, and what ensures that truth can play the substantive theoretical role assigned to it. Whilst talk about truth and correspondence might at first appear so common place as to be innocuous, you might worry that if we don't have any idea about how these

⁷ In his recent monograph Mark Jago (2018) offers a theory of truthmakers which attempts to unify the deployment of truthmakers across (1) the basis of metaphysics, (2) a theory of what truth is, and (3) their role in truthmaker semantics.

notions get into the theoretical picture, then perhaps they cannot play the robust role assigned to them. By virtue of what are certain portions of reality *truth*makers? Dyke identifies this metaphysical conception of truth, and I want to consider why it deserves the name of truth at all. I consider four possible options for what entitles Dyke to her truth-talk and find them all lacking, to greater or lesser extents, given Dyke's purposes.

First, Dyke's talk of truth might be part of some sort of austere rationalist metaphysics. Thus, perhaps we can talk about truth and correspondence without having to couch either in some generally naturalistically acceptable story, as these concepts are a priori in good standing. I take it this is certainly not in keeping with the generally naturalistic story that Dyke gives. She wants to give a metaphysic which is broadly continuous with natural science, and to thus have at the centre of her theoretical framework concepts which only have a priori rationalist justification seems deeply at odds with what else we are told. Of course, nothing that Dyke says actually rules this out, but I will take this to not capture her theoretical aims.

Second, perhaps she takes the concept of truth as her helping herself to a conceptual primitive. All theorising starts somewhere and perhaps Dyke is starting with this. If this is what she intends, this is not compelling. Whilst I am sympathetic to the thought that theories must have some minimal base of conceptual primitives, taking truth as correspondence and truthmaking to be such primitives in the context of Dyke's project is problematic. There must be some restriction on what type of primitives you are allowed given the philosophical project in which you are engaged. If Dyke intends to posit a conception of truth which cannot be motivated naturalistically, this presents a significant epistemological cost. A theoretical posit which is both taken as primitive, and which does most of the heavy lifting in your theoretical framework makes the foundations of that framework look worryingly insecure. Indeed, this seems all the more worrying vis-à-vis Dyke's particular project, which aims to sever the ties between language and metaphysics, given that the most worked out accounts of truth are generally inherited from the philosophy of language. So, in the absence of any further elucidation, her confidence that there is a language independent notion available seems unduly optimistic.

Third, she might want to endorse a Putnamian empirical conception of truth. The Putnam of *Meaning and the Moral Sciences* forwards a kind of empirical realism about truth. So we have correspondence, but an empirically satisfactory one. He claims that the primary constraint on what counts as a good truth-definition for a language is that it should be the one which best explains the behaviour of a speaker. (Putnam 1978, 41) Rationalising the behaviour of a speaker is to offer an explanation, and in order to offer this explanation we need to have an adequate handle on the connections between understanding, truth conditions, belief, action, etc. (Putnam 1978, 41) This is how correspondence gets into the picture: “‘Correspondence’ between words and sets of things... can be viewed as part of an *explanatory model* of the speakers’ collective behaviour.” (Putnam 1978, 123) Putnam thinks that the fact that there is such correspondence is what explains the success of what we can do with language, and the success of natural science. (Putnam 1978, 111)

So we get a concept of truth that is in naturalistically good standing, as it is via correspondence that we are to explain the success of human behaviour and science. Thus, the notion of truth as correspondence is secured by attending to the role it plays in this explanatory model. Could this be the sort of justification that would give Dyke the entitlement she needs? I think not. Whilst this does seem compatible with much of what she says, such a conception of truth is not robust enough to get her conception of truthmaking, along with its ontology of facts. If Dyke’s master argument is successful, it shows us that an empirically motivated truth-predicate for natural language will not yield a unique assignment of truthmakers. Whilst a Putnam-style empirical conception of truth can furnish us with a coarse-grained correspondence relation, it cannot by itself motivate the claim that the true sentences of the language are true in virtue of the existence of particular facts.

To further see why this is troubling note that Dyke thinks it is a substantive metaphysical question whether reality outstrips physical reality. (Dyke 2008, 24) For example, it is a substantive metaphysical question whether abstracta like numbers exist. If this is the case, how are we to articulate the relevant notion of success in, for example, the case of abstract ontology, without relying on representational and conceptual resources? It is not clear what the notion of

success pertaining to non-physical reality might be, without looking to language and concepts for help. You might think that the presence of mathematical terms in a regimentation of a successful physical theory licences us to infer that the objects of mathematics exist. It is because the terms of the theory correspond to reality that the theory is successful.

There are two reasons why this line of thought is incompatible with Dyke's project. First, even if we allow that the presence of mathematical language in a successful scientific theory commits us to an ontology of mathematical facts, this relation will be silent on the nature of those putative truthmakers. The empirically motivated truth-predicate can't discriminate between candidate truthmakers, and therefore cannot be the kind of relation that Dyke requires. Second, Dyke does not think that we should take the useful expressive resources of theories to tell us anything about what the world is like. But how could we come to know what abstract objects exist if it isn't in some sense via reflection on the presence of the terms which putatively name them in true theories? Unless Dyke wants to rule out the possibility that there are non-physical facts merely on the basis of the conception of truth at work (which I take it would be deeply at odds with her whole approach), rooting the notion of truth in success does not look promising.

Fourth, and finally, perhaps Dyke intends that we get to talk about truth through claiming that the principle characteristic of truth is belief. Both Lewis and Davidson think that the concept of truth has to be established as in naturalistic good standing.⁸ They do not merely help themselves to a metaphysical conception of truth as correspondence to reality. Thus, they both think it needs to be articulated in a way that gives truth an explanatory role: truth is that feature that we take sentences to have when we assent them or believe them. If we don't give truth this explanatory role, then it is not clear what entitles us to use it in our theories.

Could Dyke in principle offer a story along these lines? It is unclear. However, we might worry, as above, that such an account will not be sufficient to secure the desired ontology of facts. We might also worry that a notion of truth which

⁸ For example, Davidson (1969, 1973) and Lewis (1974a, 1975).

makes central concerns to do with belief will involve too much of a commitment to the sort of representational features of which Dyke wants to purge metaphysics. I take the central thought of Lewis's and Davidson's respective strategies to be that we only get a respectable notion of truth if it is bound to observable facts about assent. Truth is that thing that sentences are held to have when they are assented or believed. However, if the only way we can motivate a role for truth in theorising simpliciter is through this role in a theory of belief or meaning, then any role that it may play in metaphysics is going to need to be accordingly constrained by this primary function. This would then seem significantly at odds with the role that Dyke needs the concept to play. Minimally, I want to suggest that until we have a worked out theory which offers some kind of justification to underpin the work to which Dyke puts truth, it remains opaque as to why she can co-opt the notion of truth to serve her theoretical purposes.

IV. Conclusion

In this chapter I have presented a number of arguments against Dyke's philosophical project in *Metaphysics and the Representational Fallacy*. In this closing paragraph, I want to very briefly state what I take to be the important lessons to draw from the preceding discussion. First, we must be cautious that our desire to do serious metaphysics doesn't rule out the possibility of robust metaphysical projects that are anti-realist. Second, we must not allow a move towards a 'metaphysics-first' methodology to rob us of the ability to actually prosecute metaphysical theorising. Third, we must be careful with the distinction between commitments of epistemological priority and metaphysical priority. Fourth, if we propose to make use of the concept of truth in a metaphysical framework that distances itself from the role assigned to truth in the theory of meaning, then we must have a story to tell about how that concept can play the role assigned to it.

Chapter Three: Ted Sider and Truth in a Metaphysical Language

I. Introduction

The purpose of this chapter is to interrogate a language-first approach to metaphysics, where the language in question is not natural language but rather a special metaphysical language. I take as my case study Ted Sider's project in (2011) *Writing the Book of the World*. My interest in his project is that it offers a promising synthesis between approaches to metaphysics which take meaning and truth to be altogether orthogonal to metaphysical inquiry, and those which take metaphysics to be in some sense posterior to a satisfactory study of language. This method of metaphysics is therefore centrally concerned with meaning and truth, but not as these notions function in the theory of meaning for natural language. I argue that while the concept of truth plays an ineliminable role in metaphysics, both the concept itself and the assurance that it can actually play its assigned role, is inadequately articulated and motivated.

The paper proceeds as follows. First, I offer a critical outline of the main objectives and methods of Sider's project (§II). Following this overview, I take up some residual expositional matters regarding Sider's realism and the status of metaphysics (§III). In these three sections I highlight, where appropriate, the points at which Sider offers any illuminating comments on the role of truth. I finally turn to offer an evaluation of Sider's conception of the role of truth in metaphysics (§IV). I argue that there are two different ways that truth comes into the picture on Sider's view: first, via the role that truth plays in connecting the nonfundamental to the fundamental, and second, via the idea of truth as a required feature of metaphysical theory. I argue that Sider insufficiently motivates these roles for truth and thus a central part of Sider's project is worryingly opaque.

II. Sider's project

The primary claim of Sider's book is that the world has a distinguished structure, and that the job of metaphysics is to discover it. The aim of this section is to critically unpack this claim.

i. Structure and Metaphysics

Sider tells us the following:

“Discerning “structure” means discerning patterns. It means figuring out the right categories for describing the world. It means “carving reality at its joints”, to paraphrase Plato. It means inquiring into how the world fundamentally is, as opposed to how we ordinarily speak or think of it.” (Sider 2011, 1)

As means of elucidating his concept of structure, Sider draws a comparison with Lewis. He claims that for Lewis there are certain predicates that are joint-carving. These predicates pick out natural objective differences in the world. Sider tells us that structure is a generalization and extension of Lewisian naturalness. (Sider 2011, vii) He claims that “linguistically speaking, Lewis’ focus is on the predicate.” (Sider 2011, 85) For Sider, however, it is not just predicates, but quantifiers, predicate modifiers, sentential connectives, and expressions of other grammatical categories which can carve at the joints. These joint-carving notions are the fundamental ones, and Sider tells us that a fact, which I take to be an extra-linguistic entity, is fundamental when it is stated in joint-carving terms. (Sider 2011, vii) Fundamentality, then, is to be understood as a matter of structure. (Sider 2011, 5) Sider claims that the “truly central question of metaphysics is that of what is most fundamental”. (Sider 2011, 5) To prosecute answering this central question of metaphysics we will need to employ perfectly fundamental concepts. Sider contends that these

concepts will be concepts of physics, logic, and mathematics. (Sider 2011, 6) So what Sider is after is a generalised form of Lewis's approach to predicates: where "function symbols, predicate modifiers, sentence operators, variable binders, and so on, get at the world's structure". (Sider 2011, 85) Sider is clear that this doesn't mean that we treat all these other grammatical categories as predicates, but that we ask a similar kind of question. (Sider 2011, 85)

Sider's talk of conceptual fundamentality presents our first expositional challenge. Sider moves from talking about metaphysical fundamentality to talking about conceptual fundamental, without really explaining what motivates such a leap. We can see why this might be problematic. Recall the previous discussion from Chapter One on the distinction between semantic and metaphysical reduction. I claimed that there are important ways in which metaphysical and semantic primitiveness can come apart, such that a semantic primitive may not be metaphysically basic or that something metaphysically basic may be semantically analysable. Given such a possible disconnect, we should be cautious with Sider's seamless movement between talk of metaphysical and conceptual fundamentality. Is Sider employing a non-common usage for 'concept'? Is this some sort of special metaphysical notion of 'concept', where concepts are not purely representational things? Sider is not explicit about his usage of the term, however I think we should interpret his talk of concepts as still being squarely situated in realm of representation. For the notion of truth to get a grip at all, something must be a truth bearer. A more charitable interpretation of Sider's slips between talk of metaphysical and conceptual fundamentality is that it signals an important point of departure from natural language. Concepts are the tools we use to articulate our metaphysical theories and problems. However, the concepts of primary interest to the metaphysician are not those of everyday thought and talk but are rather those that belong to some distinctly metaphysical language. More on this later. Sider claims that whilst we can take the comparison with Lewis on naturalness to help illuminate what structure is, he also maintains that he cannot reductively define 'structure'. (Sider 2011, 9) Nonetheless, he thinks that a rich characterisation can be given. Sider asks us to question the extent and value of definitions; in particular, he does not think that many concepts of

philosophical interest can be reductively defined. He claims that, “we generally “understand” philosophical concepts to the extent that we know what role they play in our thinking”. (Sider 2011:9) Thus, Sider seems happy to say that we have a sufficient understanding of a concept if we know what theoretical role the concept in question is supposed to play. Sider admits that theoretical role is not all there is to meaning and reference, and as such even establishing an theoretical role for structure may not be sufficient to secure reference for the concept. However, he hopes that “this unhappy possibility is not realized”. (Sider 2011, 10) This isn’t a mere hope: Sider takes it as the purpose of the book to demonstrate how employing the concept of structure can improve our understanding of the world. He claims that “structure is a *posit*, a posit that is justified by its ability to improve our theories of these matters”. (Sider 2011, 10) He claims that posits are justified when they unify theory, and makes a comparison with unification in Newtonian physics.

I think this comparison is instructive, primarily as it reveals something about Sider’s method. Even though, as Sider acknowledges, metaphysics and physics are not the same enterprise, Sider seems happy to adopt similar epistemic practices in metaphysics as those in physics: talk of structure does important explanatory work, and as such we should believe that it is getting at something objective about the world. Note the two distinct stages here. First, although we cannot reductively define structure, we can delineate a theoretical role for it. We might therefore think of Sider as offering something like an implicit definition of structure. This shows that the concept isn’t empty of sense. Second, once we’ve got something of a handle on what the concept means, we should be content that the concept picks out some objective feature(s) of the world as the concept, so understood, can do important explanatory work. I suspect that there will have to be a particular conception of explanation at work here, in order to ensure that explanatoryness secures a referent for the concept.

Sider notes a line of resistance to the idea that we can implicitly define structure in the way he suggests is that we do not have a rich enough conception of the theoretical role of structure. How it is we get a handle on the notion of joint-carving in the case of Lewisian naturalness is via the notion of

similarity. (Sider 2011, 88) For other grammatical categories it doesn't seem we can make use of this notion. So, the thought goes, we don't really understand what it is for these categories to carve at the joints. As Sider puts it:

“As an undefined theoretical term, ‘natural’ is understood through its theoretical role. The theoretical role consists of principles specifying how naturalness relates to certain other notions, such as similarity. My term ‘structure’ has its theoretical role fixed by many of the same principles, and more besides. If the theoretical role of ‘structure’ is exhausted by its connection to similarity, and if this connection is restricted to predicates, then we have no understanding of how ‘structure’ could apply to expressions in other grammatical categories.”
(Sider 2011, 88)

In the face of such a challenge, Sider's claim is that the theoretical role of structure is not exhausted by its connection to similarity. (Sider 2011, 88) So we still have a way of understanding it, as its theoretical role is much richer. Sider contends that in addition to similarity, structure connects to many other useful theoretical notions, such as laws, explanation, metasemantics, epistemology, physical geometry, substantivity, epistemic value, etc. Such connections justify Sider's view that we can understand the concept of structure, and how it can apply to other grammatical categories aside from the predicate. In addition, Sider claims that there is a way of understanding similarity such that it is still relevant to grammatical categories other than the predicate. He thinks that in addition to talk of property similarity we can talk about fact similarity, which turns on the notion of ‘real commonality’. Sider tries to elucidate this notion of real commonality by saying that some facts ‘go together’. (Sider 2011, 89) For example, consider the facts that (1) there exists an electron, and (2) there exists a rabbit. Sider thinks that there is something in common between these two facts, a ‘real commonality’, and this is that quantification is fundamental. (Sider 2011, 89)

Let's return to the Lewis comparison. Sider considers whether, like Lewisian sparseness, talk of structure requires talk of entities. Sider's claim is that we do not need to cast talk of structure in terms of talk of entities. He claims that there is a gap between ways that we might talk and the underlying metaphysics.

For example, 'is human' gets associated with a semantic value, but my being human has nothing to do with semantic values. Sider claims that "structure facts do not concern semantic values". (Sider 2011, 91) Sider motivates this by claiming that "semantics is, like other special sciences, not fundamental." (Sider 2011, 91) He goes on:

"This is not to say that the statements of semantics are untrue, only that they are not fundamental. But if semantics isn't fundamental, the facts about carving at the joints can't fundamentally involve semantic entities." (Sider 2011, 91)

So, just because we might talk about some predicate carving at the joints, this should not be taken to reify semantic or linguistic entities. Sider therefore needs to find a way to talk about structure which does not involve commitment to entities. He suggests introducing a new locution which can be combined with expressions of any grammatical category to form new sentences. (Sider 2011, 92) Facts about structure will then be represented by this new operator, which tells us that the concept it appears in front of is a fundamental concept. For example, if φ is the new operator, then " $\varphi (\wedge)$ " tells us that the conjunction is fundamental. Sider is not concerned that we do not know of any pre-existing locution with this grammar:

"Some philosophers think we can understand only what can be defined using pre-existing resources of natural language; but this stultifying doctrine is inadequate to the evident fact of linguistic innovation within science, as well as to the initial emergence of natural language itself." (Sider 2011, 92)

This is important, and replicates the same overarching methodology we saw above. Sider thinks if we can make the novel expression "tolerably clear, and if the world contains structure corresponding to the new expression" then the introduction of the new expression is acceptable and, indeed, successful. (Sider 2011, 92) But again, this is a two-stage enterprise. First we must successfully introduce the new vocabulary, and then, if this new vocabulary does important explanatory work, we can say that it corresponds to something worldly.

What are the features of the fundamental according to Sider? He tells us that the fundamental is complete, pure, subpropositional, absolute, determinate, and fundamental. (Sider 2011, 137) How does Sider understand these commitments? He tells us that a “formulation of completeness might run as follows: *every non-fundamental truth holds in virtue of some fundamental truth.*” (Sider 2011, 105) *In-virtue-of* is not understood in terms of modality, truthmaking, or fact identity. (Sider 2011, 105) He understands purity as follows: “*fundamental truths involve only fundamental notions*”. (Sider 2011, 106) Purity leads Sider to hold that the in-virtue-of facts are not themselves fundamental. That is: facts about the relationship between the fundamental and the nonfundamental are themselves nonfundamental. (Sider 2011, 107) Sider takes structure to be subpropositional: “there are some fundamental “building blocks” – the “ultimate constituents of reality” – and the nature of reality is given by the arrangement of those building blocks.” (Sider 2011, 128) Claiming that the fundamental is determinate means something like there is no distinctive vocabulary of indeterminacy at the fundamental level and fundamental languages obey classical logic. (Sider 2011, 137) In addition, fundamentality is a fundamental notion, and some facts about fundamentality are fundamental facts. (Sider 2011, 137)

In motivating his concept of structure, Sider considers the sceptic about structure. Such a sceptic challenges the idea that we can carve reality using the ‘right concepts’ as there will be a plurality of (maybe infinitely many) ways of carving the world, and we are not in a position to say that any are better than any others. I take it this could be thought of as an epistemic claim that we aren’t in a position to know what the best ways are, or it could be thought of as a metaphysical claim that there just aren’t objectively better and worse ways of carving. What Sider needs to successfully respond to the sceptic is to make some kind of epistemically accessible distinction between those genuine features of the world, and the rest. (Sider 2011, 3) He needs to be able to make and motivate the claim that electrons are genuinely more similar to other electrons than to bison.

Sider tells us the following:

“Structure... is to be understood as objective... Whether a property, word, or concept carves at the joints has nothing to do with the place of the concept in human languages, conceptual schemes, biology, or anything like that. Thus, “fundamental” (which I use more or less interchangeably with “joint-carving” and “part of reality’s structure”) signifies a metaphysical, rather than conceptual, sort of fundamentality.” (Sider 2011, 5)

Sider contrasts what he takes to be the substantive metaphysical questions, with those that are merely conceptual. (Sider 2011, 6) The kind of answer we require for the latter sort of questions regard working out what sort of linguistic community *we* inhabit. (Sider 2011, 7) The former sorts of questions are not like this. They raise questions about whether the world has structure or not: whether the world has quantificational, modal, tensed, logical, etc. structure. Sider rightly points out that many metaphysicians talk about what is ‘really’ or ‘genuinely’ the case, and that this talk needs a satisfactory explication. Sider says that he wants the notion of genuineness to be language-independent, (Sider 2011, 4) and thus that the right way to explicate such talk of genuineness is through talk about fundamental structure. He holds that “these claims are not merely about what is *true*; they are claims about what is true at the fundamental level”. (Sider 2011, 8)

This echoes the discussion at the beginning of this section about the notion of ‘concept’ that Sider employs. Perhaps we could think of the ‘merely conceptual’ questions as those that pertain to the concepts of ordinary thought and talk, whereas the deep metaphysical concepts belong to a language of fundamental metaphysics. We are then to prosecute the project of working out what is true at the fundamental level by turning our attention to the truths stated in the fundamental language. Sider claims that a fundamental language is one whose expressions carve at the joints, and that a realism about structure commits one to a realism about fundamental languages. (Sider 2011, 8) So, in response to the sceptic who says that there is no objectively better or worse ways to carve the world, Sider responds by offering us the notion of truth in the fundamental language, where a fundamental language is one which uses joint-carving concepts. It isn’t obvious how this would satisfy the sceptic who

is worried that even if there are genuine similarities in nature we cannot know about them. I will address Sider's epistemology of structure in section three.

ii. The Method

Sider claims that the world has a privileged description. What it takes for a description to be privileged is not merely that it is true. The description must also use metaphysically perspicuous concepts, so that there is an appropriate kind of fit between the conceptual structure of a theory and reality's structure. So we want more than coarse grained correspondence. Making sense of what more is required is at times exegetically challenging. Recall, Sider holds that when I claim that some expression is joint-carving, I am not claiming that its *semantic value* is joint-carving. (Sider 2011, 90) This suggests we need to careful with talk of about correspondence, and that it should perhaps not be thought of as isomorphism. For Sider, we shouldn't explicate correspondence by taking every subsentential element of the theory to have a worldly counterpart, where this is understood in terms of *entities*.

The notion of ideology is therefore important to Sider's project. Just as much as ontology commits to 'objective content', so too does ideology. (Sider 2011, vii) He claims that a fundamental theory's ideology is just as much a part of its representational content as its ontology, "for it represents the world as having structure corresponding to its primitive expressions." (Sider 2011, viii) The pull in the other direction, i.e. to view ideology as psychological, linguistic, or conventional, Sider attributes to 'psychologizing' Quine. (Sider 2011, vii) Sider tells us:

"The success of these conceptual choices justifies belief in the existence of corresponding structures in the world. Once "ideology" has been purged of its psychological connotations, there is no barrier to recognizing a theory's logical ideology as a coequal part of that theory's portrayal of the world, and thus as being as good a candidate for carving at the joints as the theory's nonlogical vocabulary." (Sider 2011, 97)

So the ultimate aim of metaphysical inquiry is to gain insight into structure: what the world is like at the most fundamental level. But, in order to do this, we must inquire into necessity (does the best theory contain modal operators), time (does the best theory contain tense operators), ontology (does the best theory contain quantifiers), etc. (Sider 2011, 1) Given that we should not understand talk of the representational content of our ideological commitments as talk about entities (i.e. there is not some entity that corresponds to our talk of tense operators), we should understand this talk of ideological representation as the claim that reality itself is, e.g., *tensed*. So it appears that we do have a strong correspondence claim, but one that should not be understood in terms of crude isomorphism. We must therefore be very careful with some of the inflated rhetoric employed: talk of theory ‘mirroring’ reality is best viewed as metaphorical. Of course, it is still a little opaque what should replace the metaphor and I will try and bring clarity to these matters in the final section of this chapter. Talk of, for example, fundamental tense operators corresponding to tensed reality should suffice for now for getting our heads around Sider’s correspondence claims.

Sider offers an approach to metaphysics which places as central to inquiry a notion of what is fundamentally the case. This is understood in contrast to what appears to be the case. The former concerns reality and the latter concerns *our* representations of reality. Now, we may ask the question: what is the relationship between the fundamental and the nonfundamental? What is the relationship between the privileged description of reality, which only uses joint carving concepts, and our ordinary ways of thinking about the world and describing it? In what sense does the nonfundamental hold in-virtue-of the fundamental? What is the connection between them? As Sider notes, these are questions which confront any metaphysician who wants to deal in terms of fundamentality. Sider’s project, where fundamentality is understood in terms of structure, offers its own conception of the connection.

For Sider, any conception of the nature of the connection between the fundamental and the nonfundamental must hold that the connecting truths are going to be themselves nonfundamental. (Sider 2011, 110) Sider thinks that considerations from purity provide some direction for how to answer the

question of the connection. We should not be attempting to answer the question by offering a metaphysical posit: a special fundamental relationship. (Sider 2011, 110) Instead, Sider wants to introduce an undefined nonfundamental notion to explicate the connection. Although it is undefined, he thinks that it could in principle be reduced (he seems to be treating these as the same thing), but he does not think we need to do so. (Sider 2011, 112) The notion that Sider offers us is that of a *metaphysical semantics*. (Sider 2011, 112) This is a close relative of the notion of ground, but Sider claims he offers a slightly different linguistic approach. (Sider 2011, 112)

Sider claims that a metaphysical semantics is a semantic theory with two distinctive features. First, “meanings are to be given in joint-carving terms”. (Sider 2011, 112) So, if the metaphysical semanticist offers a truth theory, the right-hand side of the truth conditions must be given in purely joint carving terms.¹ Second, its explanatory goals differ from linguistic semantics. (Sider 2011, 112) On the one hand, metaphysical semantics seeks to explain how what we say fits with fundamental reality. On the other, it is not constrained by the sort of cognitive or syntactic data points that constrain linguistic semantic theory. (Sider 2011, 112-113) The metaphysical semanticist is “free to assign semantic values that competent speakers would be incapable of recognising as such, for she is not trying to explain what a competent speaker knows when she understands her language”. (Sider 2011, 113)

There is an interesting question here about precisely what it is that a competent speaker could not recognise, and quite how this differs from natural language semantics. I take it the thought is that when assigning semantic values in the metaphysical semantics there is no requirement that such assignments be cognitively transparent. Of course, if we are talking about semantic value assignment in natural language it is far from clear that such assignment should be cognitively transparent. Perhaps the assignment of truth conditions needs to be in some way cognitively constrained, e.g. if I understand the sentence I should be able to recognise the truth conditions. But, in claiming there is

¹ Sider is open to the thought that there are other options for giving the meaning which don't proceed truth theoretically, e.g. assertion-conditions etc. but claims that for factualist discourse, we should prefer a truth-theoretic approach. (Sider 2011, 113-114)

something significantly different going on in the metaphysical semantics from natural language semantics, such that the former is less constrained, I worry that Sider may be conflating the theory of sense and that of reference. If the claim is that in metaphysical semantics a competent speaker couldn't recognise the extension of a term, then this seems innocuous and continuous with the natural language semantical project. If the claim about recognition is something more than this, it is not clear quite what it is supposed to be.

Putting these concerns about cognitive transparency aside for a moment, with the shape of the project identified, Sider questions the degree to which we can expect ourselves to be able to actually offer metaphysical truth conditions.

(Sider 2011, 116) In considering this, Sider makes some general claims about the failures of projects of conceptual analysis. (Sider 2011, 117) He thinks that the problems with conceptual analysis are exacerbated when we try to give truth-conditions in joint-carving terms. But Sider thinks we can construct a 'toy' metaphysical truth-conditions, for the purposes of getting a handle on what sort of shape they will have, without committing us to getting them to fit perfectly with usage and in purely fundamental terms. (Sider 2011, 117) The purpose of the toy metaphysical semantics "is not to *be* a real metaphysical semantics, but rather to convince us that *there is* a real metaphysical semantics, even if that metaphysical semantics is too complex for us to discover". (Sider 2011, 117)

To close this section, I want to register that I find this stage of the dialectic challenging. Sider seems to be committed to the view that a central and indispensable part of his project may be impossible for *us* to execute. I take it he thinks that in principle such a project is possible, but maybe not for creatures like us. Sider has told us that the metaphysical semanticist is "free to assign semantic values that competent speakers would be incapable of recognising as such, for she is not trying to explain what a competent speaker knows when she understands her language". (Sider 2011, 113) Whilst it does not seem that it should be a requirement that the semantic *values* be recognisable as such, Sider's claim seems to be stronger than that: that, perhaps, we cannot even recognise the truth conditions. If the truth conditions

we need to assign to execute the project are inaccessible, then large parts of metaphysics will likewise be inaccessible.

Now, perhaps this will not worry Sider. He is a committed realist (as I will discuss in the next section), so the fact that reality is going to outstrip our ability to theorise about it is perhaps not something he sees as a problem. But it is dialectically curious that Sider is arguing for a way of doing metaphysics which may be inaccessible to us. What I take this to highlight is the contingency of cognitive success when we attempt to offer theories about what the world is like. There are a great many philosophical and scientific projects which we want to be able to execute but, given our cognitive limitations, we may not be able to do so. Some might take our inability to execute such projects as demonstrating that there is something problematic with the project itself. Perhaps you might think that if metaphysical truth conditions are ineffable, then this points us in the direction of the impossibility of metaphysics. Here, again, Sider deploys the analogy with science. We are cognitively limited in physics, but we don't generally take this to show that there is something wrong with the methods of physics.² If you are happy to buy both realism and the contingency of cognitive success in physics, but you don't want to accept that package deal in metaphysics, then it needs to be shown that metaphysics is by its nature worse off than physics. I will now turn to consider such questions about the status of metaphysics.

III. The Status of Metaphysics

In this section there are three residual expositional matters which I want to discuss, all of which pertain to Sider's view of the status of metaphysics. First: his view on the epistemology of metaphysics. Second: his 'knee-jerk realism'.

² Of course, it must be noted that a number of prominent philosophers of science have defended the view that our epistemic limitations with respect to fundamental physics should indeed impact what our metaphysical commitments should be. See, for example, van Fraassen (1980).

Third: his claim that something like his conception of metaphysics is unavoidable.

iii. The Epistemology of Structure

Sider urges that we adopt the Quinean view of ontological commitment and apply (something like) it to the epistemology of structure. (Sider 2011, 12)

“As a general epistemology of metaphysics I prefer the vague, vaguely Quinean, thought that metaphysics is continuous with science. We employ the same criteria... for theory choice in metaphysics that we employ outside of metaphysics.” (Sider 2011, 12)

We start by finding the best theory, which is the one that is most parsimonious, explanatory, coherent with other good theories, etc. Whilst we may have to accept that these criteria give us less guidance in metaphysics than they do in science, Sider thinks this follows from the speculative and uncertain nature of the discipline, and doesn't seem worried. (Sider 2011, 12) Sider claims that the Quinean holds that we then “believe the ontology of your best theory”. (Sider 2011, 12) Sider motivates this by noting that the ontological commitment is a part of the theoretical success of the good theory. Sider then urges that we extend such a picture to include a non-psychologised view of ideology. So, we should believe the ideology of the best theory; in particular, we should take it to carve at the joints. (Sider 2011, 12) If you buy this story then, Sider thinks, we can have evidence for our claims about structure. The story about the epistemology of structure, however, goes beyond this vague Quinean picture. Sider thinks that when we search for the best theory we search “simultaneously for a set of concepts and a theory stated in terms of those concepts”. (Sider 2011, 13) So, metaphysical or empirical inquiry is about finding the best theory (according to the identified theoretical virtues) *and* about finding the best concepts to articulate that theory.

Sider claims that this epistemological story provides further vindication for the concept of structure itself. He claims:

“To evaluate proposed trades of ontology for ideology – which we must do in order to choose which fundamental theory to believe – we need to speak of joint-carving for expressions other than predicates.”
(Sider 2011, 87)

So, the thought goes, if we do not extend our conception of structure beyond the predicate, we lose one of the most significant methods of adjudicating between competing metaphysical theories. I take the proposed insight here to be that part of the Quinean epistemology involves some kind of cost-benefit analysis where we weigh up competing articulations of theory. In order to do this we need to get a handle on which expressions are ideologically committing, and Sider’s framework purports to offer a method for doing this.

We might worry, however, whether we actually need Sider’s special notion of structure to do this. To help to illustrate how we might resist this, it is worth reflecting on the degree to which Sider’s own proposal is any different from Quine’s. Although Quine is most explicit about offering an explication of ontology, he arguably intends this to be extended, *mutatis mutandis*, to other areas of classical metaphysics. Following Divers (2017), we can understand Quine’s modal anti-realism as deriving from his view that Best Theory does not contain any modal operators or connectives. Unlike respectable ontology, where we are concerned with quantifiers and variables, respectable modal metaphysics concerns itself with the question of whether quantified modal logic deserves a place in Best Theory. It is hard to see how in the details this Quinean research programme differs from Sider’s own. In both cases, questions of metaphysical interest are to be answered by attending to the logical constituents of the metaphysically salient discourse.

I suspect the difference is not so much in the proposed methodology as such, but in the view taken of the alethic status of the subsequent theories. For Sider, the claim that best theory contains some ideological commitment is accompanied with a strong claim about how the structure of theory corresponds to the structure of the world. If best theory contains tense operators, then this reflects the tensed nature of reality. Truth for Sider, at least the type the metaphysician is interested in, has to do with a matching between the structure of the best theory, and the structure of the world. Quine, by

comparison, is not going to be happy with such a substantial, subsentential notion of correspondence given his view of the indeterminacy of reference.³ Thus, although Sider thinks that his Quinean epistemological story provides motivation for the idea of structure itself, reflecting directly on Quine's method seems to call into question any such motivation. To get Sider's required motivation for the new theoretical notion of structure, we need a further commitment to a strong correspondence conception of truth.

Sider considers an epistemological challenge to his concept of structure which follows from the primitivism he endorses. (Sider 2011, 11) The challenge proceeds as follows: how can we know about structure if we cannot get a handle on it in more familiar terms. This challenges our very understanding of the concept for if we cannot know any facts about structure, then we cannot isolate a theoretical role for it (as it will be cut off from our other concepts). Given Sider's view of how theoretical terms get their meaning, this threatens our understanding of the concept itself. At this point, Sider encourages a fallibilist epistemology. (Sider 2011, 12) He concedes that this move may be resisted by those who think that the subject matter of metaphysics is so far removed from our normal fallible epistemic practices. But, to abate this concern, he suggests a comparison with mathematics, logic, and particle physics, where he claims that in these areas too our contact with them is unlike the world of middle sized dry goods. (Sider 2011, 12)

Sider admits that his epistemological story is "superficial and birds-eye", but claims this is not merely a problem with structure but a wider problem. (Sider 2011, 15) Sider claims that many metaphysicians view their work as not being conceptual archaeology, but rather as being more akin to speculative science. (Sider 2011, 72) He seems to think that although epistemological concerns are serious matters to be considered, the metaphysician does not need to be any more worried than the mathematician or the particle physicist. As I suggested at the end of the last section, we might again question whether metaphysics is really any worse off than these other theoretical disciplines. In this context, at least, I think we have some reason to be sceptical of the comparison. It is not

³ For Quine's discussion of the indeterminacy of reference see Chapter Two of (1960).

clear to me that either of these other theoretical enterprises by themselves are committed to the sort of fundamental metaphysics Sider endorses. Indeed, whether they are such enterprises seems to be precisely what philosophers of physics and philosophers of mathematics might argue about. To assume that these are joint-carving enterprises, from which we can borrow epistemic practices, seems to put the cart before the horse.

Of course, Sider is fully willing to concede that metaphysics is, in important respects, dissimilar from these other disciplines:

“Some of the metaphysically deepest disputes... are conceptually shallow in that they have few implications outside rarefied metaphysics.” (Sider 2011, 74)

He acknowledges that this might give one reason to think that there is something odd going on in metaphysical debates. Nonetheless, he is not worried about it, and is happy to maintain the comparison between metaphysics and other theoretical pursuits. (Sider 2011, 74) However, you might think that the above quotation shows that there is a more serious disanalogy. Think of a comparison: interpretations of quantum mechanics may be rarefied but particular views have consequences (even if it is not always clear what the formalism says about what the world is like). The fact that Sider acknowledges that some of the metaphysically deepest disputes may have little to no consequences outside of metaphysics may give us a reason to resist adopting the fallibilist epistemology he recommends, as it isn't immediately obvious what would count as 'getting it wrong'.

iv. Knee-Jerk Realism

This brings me on to the next topic for discussion: Sider's realism. Sider tells us the following:

“A certain “knee-jerk realism” is an unargued presupposition in this book. Knee-jerk realism is a vague picture rather than a precise thesis. According to the picture, the point of human inquiry – or a very large

chunk of it anyway, the chunk that includes physics – is to *conform* itself to the world, rather than to *make* the world. The world is “out there”, and our job is the wrap our minds around it.” (Sider 2011, 18)

An interesting thought to entertain regarding this admittedly vague formulation of realism, is that it is not clear that this way of thinking about realism can get Sider quite what he wants: knee-jerk realism is seemingly compatible with other more deflationary conceptions of metaphysical inquiry. For example, Amie Thomasson explicitly argues against Sider’s joint-carving project, but she still thinks that the world is “out there”. (Thomasson 2015, 296-317) Commitment to knee-jerk realism won’t alone get you realism about joint-carving. Nonetheless, Sider suggests that a consequence of knee-jerk realism is that “the physical description of reality be objectively privileged. And a natural account of the privilege is that physical notions carve at the joints.” (Sider 2011:20) It is not entirely clear what motivates this jump. Sider acknowledges that going realist in this way means the epistemological question then confronts us ‘head on’. (Sider 2011, 42) However, Sider further claims that we have a standard way of addressing such problems as underdetermination of theory by observation: we use our theoretical virtues to adjudicate. I take it then, that Sider’s brand of realism and his epistemological story are very tightly intertwined.⁴

v. Unavoidable Metaphysics

The final point to discuss in this section is what Sider thinks about the unavoidable nature of metaphysics. For Sider, metametaphysical critiques are metaphysical in nature. (Sider 2011, 82) I take the thought to be that in order to forward some kind of anti-metaphysical thesis, we need to claim that the debate is in some sense nonsubstantive. For Sider, to say of a question that it is nonsubstantive, is to say that there are different answers to the question each

⁴ Of course, this invites similar questions to the previous section about the relationship between Quine’s and Sider’s project. Even equipped with his preferred epistemological story, this still doesn’t seem to be enough to get Sider quite the brand of realism he wants.

of which come out true on different interpretations of the salient terms, and no candidate interpretation carves at the joint better than any other. (Sider 2011, 74) Thus the machinery Sider gives us to assess nonsubstantivity is whether or not a theory is cast in joint-carving terms. So, Sider claims, merely reflecting on language and reasoning cannot tell us whether a debate is nonsubstantive as we actually need to know something about the structure of the world in order to claim that a debate is nonsubstantive. He claims:

“There is no ametaphysical Archimedean point from which to advance deflationary metametaphysics, since any such metametaphysics is committed to at least this much substantive metaphysics: reality lacks a certain sort of structure.” (Sider 2011, vii)

Sider claims that some of the challenges to metaphysics as an enterprise rest on oversimplified conceptions of language or knowledge, and, since they threaten the science of unobservables as much as metaphysics, he is not worried about them. (Sider 2011, 67) What he takes to be a more serious threat to metaphysics comes from what he terms metaphysical deflationism, which holds that answers to specific metaphysical questions comes from reflecting on linguistic or conceptual rules, rather than facts about “the world”. (Sider 2011, 67) It isn’t clear that facts about linguistic rules aren’t about the world, but I take it that Sider’s claim is rather that facts about linguistics are not fundamental. The particular brand of deflationism that Sider seems to have in mind is one which holds that for a metaphysically contested sentence, there are multiple equally good ways of interpreting the sentence such that we can all agree that the sentence is true. (Sider 2011, 68) Sider suggests the following as a rough summary of the go-to move for this type of deflationist: “there’s a perfectly good way to talk under which S is clearly true”. (Sider 2011, 68)

Sider notes that the guiding thought behind this move is that meaning is largely determined by use. (Sider 2011, 68) So the thought is that there are going to be a range of candidate interpretations of the sentence, and the sentence will come out true under each interpretation, but that usage does not settle which of the interpretations is correct, so the question itself is merely conceptual. (Sider 2011, 69) Sider claims that in the case of fundamental metaphysics “the most straightforward way to resist deflationism is to claim that the crucial

expressions in the debate carve perfectly at the joints”. (Sider 2011, 71) Here again, Sider reinforces the idea that the serious metaphysician is not merely after truth, but after truth in the right terms. This of course assumes that such a metaphysician can help themselves to structure in the first place. However, Sider seems to think that if the deflationist has no conception of structure at all then, like the attacks on metaphysics above, they are not in a position to make a distinction between metaphysics and the science of unobservables. (Sider 2011, 69) Consequently, we needn’t worry about such critiques.

The claim that Sider makes is that there is no way out of metaphysics; there is no ametaphysical high ground, where we can avoid making distinctly metaphysical claims. Sider tells us: we cannot avoid metaphysics. (Sider 2011, 83) This is not to say that a deflationary position will never be attractive, but rather that that is just more grist to the metaphysical mill. If the various candidate meanings are all just as good as each other, then the world lacks a certain kind of structure. If it didn’t lack this kind of structure, the terms would pick out the joint-carving meanings. In order for the deflationist to make their case they need to deny some kind of structure. So the options for the anti-metaphysician are, according to Sider, (1) to deny that there is any sensible notion of structure at all, but in so doing preclude a robust view of the science of unobservables, or (2) to deny that some particular kind of structure exists, but this is still to make a substantive claim about the nature of reality.

This conception of the options for the anti-metaphysician seems to me to be not exhaustive. More precisely, Sider implicitly presents the options of *endorsing* some commitment and *rejecting* that commitment as being exhaustive. But it seems like there is another attitude we can take to that commitment: I can refuse to assert it without denying it. This ‘anti-metaphysics’ position is epistemic in nature, and opens the possibility of adopting a quietism about particular metaphysical debates. It is unfortunate that Sider doesn’t consider such a position as it seems like a simple example of how metaphysics can be avoided.

IV. Language Truth and Metaphysics

With a comprehensive picture of Sider's project on the table, in this section I want to try to elucidate Sider's views on the role of truth in metaphysics. Much of what Sider says about these matters are suggestive, so I will be engaged in quite a constructive exegesis. There are two different ways that truth comes into the picture on Sider's view. First there is the role that truth plays in connecting the nonfundamental to the fundamental. Second, there is truth as a required feature of metaphysical theory. Both of these deployments of the concept trade in a conception of truth as correspondence to reality. In this section I want to evaluate what motivates both of these deployments and say something about the relationship between them.

vi. Metaphysical and Linguistic Analysis

In the context of his discussion of toy metaphysical truth conditions, Sider comments on the failures of projects of conceptual analysis. He offers the following comment:

“Conceptual analysis is out of fashion in metaphysics, but there is uncertainty about what its replacement should be. Reduction? Supervenience? Realisation? The proposed replacements have tended to be either inadequate or unilluminating (or both). The recent trend is to think in terms of a kind of purely metaphysical analysis. There's a bad idea in here mixed in with good one. The bad idea is that we should posit a fundamental gizmo for the relationship between analysans and analysandum (truthmaking, necessity, fact identity, ground). The good idea is Armstrongian: metaphysical analysis is not linguistic analysis. But that leaves a big gap: if not linguistic analysis

then what? I say: metaphysical truth-conditions – toy models of them at any rate.” (Sider 2011, 118)

Sider notes the tradition in the philosophy of language to view metaphysical and linguistic inquiry as tightly constraining each other, and points to Davidson as an example of a philosopher who holds that metaphysics is constrained by linguistic theory. In particular, a good semantic theory should where possible take sentences to be true. Sider suggests that the order of priority should run the other way: metaphysics should constrain the philosophy of language. (Sider 2011, 122) Metaphysical arguments can have a serious impact on how we semantically analyse. Sider claims that one advantage of viewing things this way is that there can be a looser connection between metaphysics and semantics than on Davidson’s account. So, the linguistic semantics can take ordinary sentences as true but this needn’t ontologically commit them, as ontology concerns fundamental existence. (Sider 2011, 122) Sider is open to the thought that human thought and talk can influence metaphysical theorising, but only inasmuch as this is because they are real phenomena and so must fit into the final metaphysical picture. What Sider does not want is a mere “reading off one’s fundamental metaphysics directly from the structure of thought and talk”. (Sider 2011, 123)

We should pause here to think about Sider’s challenge to Davidson. In particular, we should question what motivates those who think that metaphysics is going to be constrained by linguistic theory. I take the basic and central commitment here to be that semantic concepts only earn their keep if they play some role in empirical theories, e.g. of behaviour, mind, sociality. This requires that there have to be discernible connections between the application of these concepts and observable facts. So, in the case of truth, truth earns its keep in an empirical theory of communication. Adding the concept of truth to our conceptual repertoire makes an explanatory difference when we are trying to give empirical theories of understanding, communication, belief, and interpretation. In particular, we only get a respectable notion of truth if it is bound to observable facts about assent. Truth is that thing that sentences are held to have when they are asserted or believed. If the only way we can motivate a role for truth in theorising

simpliciter is through its role in the theory of meaning, then any role that it may play in metaphysics is going to need to be accordingly constrained by this primary function. If we cut truth off from the linguistic analysis, we render it undermotivated at best or unintelligible at worst. If this is the way you are thinking about the concept of truth, it is not implausible to hold that switching the order of priority would bankrupt truth as an important concept for metaphysics.

Another consideration which should make us question the switching of the order of priority that Sider recommends here, is that he doesn't satisfactorily disambiguate the different ways in which linguistic theory may have claim to priority. As we saw in the last chapter, there is an important distinction to be mindful of when discussing metaphysical method and priority. Just because I may have to methodologically proceed from language to world, this does not commit me to the view that the language is in anyway metaphysically prior to the world.

A plausible explanation for why Sider rejects the Davidsonian project is that Sider is committed to a very different conception of natural language semantics. Sider claims that adopting his view of structure can help to answer the radical semantic sceptic. He makes two important claims. First that reference relations are explanatory relations, and second that explanatory relations are joint-carving ones. He takes these two views together to defend a reference magnetism response to permutation arguments for radical semantic indeterminacy. The thought being that, as reference is an explanatory relation, a view of reference determination which allows for crazy permuted models is an explanatorily useless view, and as such should be rejected in favour of one which gives reference the right kind of explanatory role. (Sider 2011, 28) He claims that "reference must have the right sort of basis in the fundamental if it's to be explanatory." (Sider 2011, 29) Sider concedes that 'the right sort of basis' requires clarification, but thinks we can see in a clear cut way when the wrong sort of basis is present. (Sider 2011, 29)

So, reference magnetism answers the radical semantic sceptic. This is because reference is an explanatory relation and explanatory theories must be stated in joint-carving terms. We therefore need structure to answer the radical semantic

sceptic. Sider claims that a “more reasonable metasemantics will allow a role for a non-observational and externalist determinate of meaning: the world’s structure.” (Sider 2011, 42) He claims we can thus reject empiricist conceptions of meaning. This perhaps explains why Sider does not helpfully engage with *why* certain kinds of projects think that the theory of meaning constrains metaphysics. Sider is theoretically very far away from such empirical semantic projects, and thus does not appear to try to understand why, starting from such a position, you might think that linguistic analysis is tightly bound to metaphysics. That is: because of the role afforded to truth in both enterprises. Sider’s claims about the relationship between metaphysics and linguistic analysis extend further than his claim that metaphysics may depart from projects of linguistic analysis. He doesn’t merely think that we shouldn’t read off our metaphysics from the structure of thought and talk, Sider suggests that often we should not conduct metaphysics in ordinary language at all. As the ordinary language expressions may not carve at the joints, such theorising would lead to nonsubstantive debates. In such cases he suggests the following:

“But we could discard \exists , and enter the metaphysics room, so to speak. We could replace the ordinary expression \exists with an improved expression \exists^* that we stipulate is to stand for the joint-carving meaning in the vicinity. The question we ask in the metaphysics room, cast in terms of \exists^* rather than \exists , is substantive. Indeed it is *superior* to the original question, for it concerns reality’s fundamental structure, rather than its merely conventional or projected aspects.” (Sider 2011, 74)

Sider is appropriately sensitive to the fact that even though we have moved into a ‘superior’ language for theorising, we need to be able to tell a story about how the new terms of the discourse derive their meanings. In ordinary language we have a story, perhaps something about use, so when we leave ordinary language we need an alternative:

“Abandoning ordinary language is indeed often a bad idea, but when it is, that is because there is no *other* way to anchor the debate, no other way to explain the meanings of the crucial words without trivialising

the debate. But joint-carving meanings give us another anchor.” (Sider 2011, 75)

So, joint-carving meanings are to give us our required anchor. Sider provides us three conditions which must be met in order to legitimately introduce a new term. Take a new term ‘Cause*’. This new term “should stand for a meaning which i) carves at the joints (perfectly); ii) fits our use of ‘cause’ better than any other joint-carving meaning; and iii) fits our use of ‘cause’ well enough”. (Sider 2011, 76) Sider considers the thought that such successful stipulations are not possible; perhaps because reference magmatism is false or some non-referential account of meaning is true. Sider thinks that if this were the case then these views must nonetheless allow for something like successful stipulation to be true in science or else this would be a reductio of the view. (Sider 2011, 76) Given that successful stipulation is allowed in science, Sider thinks the case will be parallel for metaphysics.

At this point we might wonder what the relationship is between the already discussed metaphysical semantics and Ontologese (the name he gives to our new and improved language for metaphysical theorising). In Ontologese, all expressions are joint-carving. Likewise, the metalanguage for the metaphysical semantics needs to be perfectly joint carving. I therefore take the best interpretation of the relationship to be that Ontologese provides the metalanguage for the metaphysical semantics. When offering some metaphysical truth condition for a sentence pertaining to the realm of appearance, the right-hand side of this truth condition is to be stated in Ontologese. It may help to work through an example. Take a toy metaphysical truth condition like the following:

‘x’ is true iff x*

The sentence on the left-hand side stands for an ordinary derivative representation. The right-hand side stands for a sentence of Ontologese. So if we take the left-hand side to be the sentence ‘there used to be dinosaurs’, then, if tense is fundamental, the sentence of Ontologese will contain a tense operator. Further, truth plays an important explanatory role here in connecting the nonfundamental to the fundamental.

There is an interesting tension between the different sets of demands placed on the metaphysical semantics from a natural language theory of meaning, a tension which comes up many times in Sider's book. On the one hand, given what Sider says about how the introduction of new terms in Ontologese is constrained to some extent by ordinary usage, then the metalanguage of the metaphysical semantics is going to be accordingly constrained. Of course, his magnetism picture of reference assignment suggests that there are going to be competing metasemantic constraints: naturalness and ordinary usage. But it still seems of note that to some extent at least the articulation of metaphysical truth conditions is constrained by ordinary usage. Sider also claims that "a metaphysical semantics must successfully explain the linguistic behaviour of the population in question. (Sider 2011, 114)

On the other hand, however, Sider is explicit that it is not important for metaphysical theory and linguistic theory to "neatly mesh". (Sider 2011, viii) I take the thought to be that the fundamental metaphysics underlying a discourse might have a structure that is very different from that suggested by the discourse, so sometimes our introduced terms will not have a close cousin in ordinary English. Another comment that Sider makes on the separation of natural language semantics from metaphysical semantics is that one traditional problem for semantics that the metaphysical semanticist need not concern herself with is that of Frege's puzzles about belief. (Sider 2011, 11)

It is initially challenging to see how these two sets of commitments fit together. The metaphysical semantics is such that it is both constrained by existing language use and semantic knowledge, but also it does not have to be able to explain certain semantic puzzles or perfectly map into natural language semantics. I suspect that some of the previous discussion about the requirements of cognitive transparency for the metaphysical semantics may help here. In what sense does the metaphysical semantics need to explain linguistic behaviour? Sider is not maximally perspicuous about what he intends, so I suggest that the best way to think about it is as the claim that although the metaphysical truth conditions need not be cognitively transparent, they are nonetheless constrained by established patterns of dissent and assent.

vii. The Role of Truth in Metaphysics

We saw that one of the central mechanisms of Sider's fundamentality story is that of metaphysical truth conditions, and in the previous section I attempted to elucidate this notion, and explain how it relates to linguistic analysis. We saw how truth plays a role in connecting the nonfundamental with the fundamental. In this section I want to unpack quite what this concept of truth is, and what enables it to play the role that it does in Sider's framework.

Although Sider does not dedicate much space to the explicit discussion of truth, when he does talk about it, it is often in terms of the truth of a whole theory. I take it that the relationship between these two deployments of the concept is the following: the metaphysical truth conditions which are satisfied collectively compose the metaphysical theory.

When introducing the notion of structure in the first section, I drew attention to Sider's claim that "for a representation to be fully successful, truth is not enough; the representation must also use the right concepts, so that its conceptual structure matches reality's structure." (Sider 2011, vii)⁵ We are to understand the 'betterness' of some representations over others not in terms of truth, which we take for granted that we want from a representation, but in terms of joint-carving terms. (Sider 2011, 19) Sider's most salient and repeated claim about the role that truth plays in metaphysics is the following:

"What we care about is truth in joint-carving terms, not just truth."
(Sider 2011, 63)

What should we make of this claim? Sider suggests that good scientific theories must be cast in joint-carving terms. He suggests that a useful way of understanding this is in terms of explanation: "theories based on bizarre, non-joint-carving classifications are unexplanatory *even when true*." (my italics, Sider 2011, 23)⁶ What I take to be salient about this is that truth just isn't the primary

⁵ Of course, I previously cautioned we need to be wary of this talk of mirroring or matching.

⁶ There is another salient question here about quite what role explanation is playing in Sider's framework. We might wonder why something being explanatory carries any deep metaphysical significance. It seems that the important metaphysical question is what the telling discourse is.

notion of concern in metaphysics; truth is still required, but we need more than that. One interesting thing about this is that it takes the concept of truth itself as a given, and then builds on top of it.

However, we might wonder why Sider is permitted to do this. What gets us the concept of truth in the first place? I conjecture that the most natural way to think about Sider's conception of truth is something akin to Putnam's suggestion in the early papers of *Meaning and the Moral Sciences*. We have a fairly minimal, Tarskian view of the truth predicate, which we then situate in a story about science and success. We then get a picture where truth has the formal properties we want, but is fit to play a role in realist scientific theory and explanation. So we have an empirical conception of truth and then we project this notion we get from the empirical project onto this metaphysical one. Sider certainly takes truth to be an explanatory relation, however, he is not explicit about whether it is explanatory in this Putnamian sense. Nonetheless, I think such a commitment is in the background. It seems to me the most expedient way to get him a scientifically respectable notion of correspondence.

From here, we should now question how should we understand Sider's talk that 'truth isn't enough'? It is clear that in addition to truth Sider also wants joint-carving. It seems like there could be two different ways of understanding this dual requirement. First, perhaps there is some kind of truth pluralism in the background: there are multiple concepts of truth (or perhaps one concept, but multiple properties) and Sider is then specifying that the concept/property he wants is a joint-carving one. Second, perhaps there is a fairly minimal conception of truth, and then Sider wants to place certain joint-carving restrictions on the statement of the truth conditions. However, we might worry that this kind of requirement on truth does not sit well with the Putnamian empirical conception of truth I previously suggested. It doesn't *prima facie* seem like Putnam's arguments in *Meaning and the Moral Sciences* will motivate the fine-

Sider's answer seems to be that the most *explanatory* discourse is the telling discourse. What does Sider have to say to someone who maintains that explanation has to do with mode of presentation, not worldly stuff? For some interesting discussion of the non-objectivity of explanation, see Thompson (2016).

grained correspondence relation that Sider wants, and the kind of metaphysical restrictions he wants to place on the specification of the truth conditions.

Perhaps then my suggestion that we should understand Sider's conception of truth in (early) Putnamian terms, was too quick. However, given that he does not offer any comment on how he is thinking about it himself, and if the Putnamian view is not up to the theoretical task that Sider assigns it, what are his alternatives?

The most promising place to look for an alternative conception of truth which is compatible with placing such restriction on the specification of the right-hand side of truth condition comes in his discussion of substantivity. This contains his most explicit discussion of truth. Sider makes a case for understanding notions of substantivity, conventionality, and subjectivity in terms of joint carving. He claims that understanding the nature of the connections between these concepts can be understood as uniting under a single thesis about epistemic value: "it's better to think and speak in joint-carving terms". (Sider 2011, 61) The reason for this rests on a very general claim about the nature of inquiry. Sider claims that the goal of inquiry isn't simply to believe true things, but rather to have the content of one's beliefs reflect the structure of the world. He claims that "wielders of non-joint-carving concepts are worse inquirers". (Sider 2011, 61) This isn't merely a claim about the practical value of joint-carving beliefs; Sider thinks that it is a constitutive aim of belief, along with truth. (Sider 2011, 61) That belief has this aim does not have to be acknowledged by all the members of an epistemic community, but it is the standard by which belief should be evaluated nonetheless. (Sider 2011, 61)

What reason do we have to posit this extra norm of belief? Sider gives two reasons. First, he claims that, like truth, beliefs aim to conform to the world. (Sider 2011, 62) Second, Sider takes scientific inquiry as a paradigm of good inquiry because it is conducted in joint-carving terms. (Sider 2011, 61) Sider motivates these two claims by considering the alternative. He claims that:

"If there is no sense in which the physical truths are objectively better than the scrambled truths, beyond the fact that they are propositions

that we have happened to have expressed, then the postmodernist forces of darkness have won.” (Sider 2011, 65)

Rhetoric aside, it seems that Sider’s thought here is that if we only have truth, but not the special joint-carving kind, we may as well give up on the whole metaphysical enterprise (or perhaps inquiry tout court). That is, if we have truth but no joint-carving, this “diminishes the value of truth.” (Sider 2011, 62)

We might wonder what to make of this. Sider doesn’t really tell us anything about the concept of truth, apart from the fact that it is correspondence. If you started out sceptical of philosophers helping themselves to a conception of truth as correspondence to reality, Sider doesn’t provide anything to assuage those concerns. What he does try and do is motivate the idea of truth in joint carving terms. He does this by grappling with some fundamental issues about the general purpose of inquiry. In order for a metaphysical theory to be legitimate it must be both true and articulated in joint carving terms. The reason for this is because the epistemic, and perhaps practical, value of theories which are both true and joint carving are greater than any rivals. However, if I’m thinking about the demands of inquiry and about truth, I’d be inclined to think that what grounds the epistemic and practical value of such theories is that they increase success. I’ve already suggested that a Putnamian empirical conception of truth won’t get you Sider’s joint carving restriction on truth, and it isn’t obvious that there is a notion which could satisfy all the demands he places on it.

It should be clear that truth does play an important role for Sider, but that not any old conception of truth will do. I have claimed that it is a conception of truth where there are restrictions placed on the specification of the truth conditions. The theoretical task is to then articulate what motivates this conception. I conclude this chapter claiming that Sider has not adequately met that task. Truth is required to play a role in Sider’s project that it is not obvious that it can play. And this should be worrying, given the central role assigned to truth via the metaphysical semantics. Sider dedicates a lot of time and care to motivate the concept of structure. However, the lack of similar motivation for the concept of truth renders his approach incomplete. Where does this leave us? Either with a central part of Sider’s picture remaining worryingly opaque

and unmotivated, or perhaps with some kind of background commitment to a rationalist conception of truth which, without argument, seems equally unmotivated.

V. Conclusion

In this chapter I have shown that Sider presents a robustly realist conception of metaphysics, where the epistemological status of the enterprise of metaphysics should be seen as on par with that of natural science. I have argued that truth play a significant role in Sider's proposed method for metaphysics, first, via the role that truth plays in connecting the nonfundamental to the fundamental, and second, via the idea of truth as a required feature of metaphysical theory. I concluded that Sider insufficiently motivates these roles for truth and thus a central part of Sider's project is worryingly opaque. In a similar spirit to the conclusion of the previous chapter, the important take away from this chapter is that if we propose to make use of the concept of truth in a metaphysical framework then we must have a story to tell about what legitimises this use.

Chapter Four: Donald Davidson, and Semantic and Ontological Indeterminacy

I. Introduction

The purpose of this chapter is to critically evaluate Donald Davidson's conception of the role of truth in metaphysics, and the status of the consequent metaphysical inquiry.¹ Davidson's project presents an interesting case study. A semantic conception of truth is his central theoretical notion, which earns its keep through the role it plays in the theory of meaning. Metaphysical commitments are generated when the true sentences of the language are analysed to reveal quantificational structure. Davidson's project bears many similarities to Lewis's project: a necessary prerequisite to metaphysics is the analysis of language. However, unlike Lewis, Davidson is not committed to the view that there is subsentential semantic determinacy, and this has a rather radical impact upon the status of the metaphysical enterprise.

The chapter proceeds as follows. First, I outline Davidson's theory of meaning (§II). This has two distinct components. I address the role of truth in the theory of meaning, and the project of radical interpretation, focussing on how the methodological presuppositions of radical interpretation introduce some notion of correspondence. Second, I construct a picture of Davidson's metaontological views (§III). Again, this has two dimensions. I present his explicit views on the proper methodology of metaphysics, and then I show how Davidson, like Quine, is committed to ontological relativity. Third, given the previous discussion, I evaluate the conception of metaphysics with which Davidson's project leaves us (§IV). As a means of evaluating Davidson's

¹ The discussion of Davidson in this chapter is restricted to his early (pre-1980s) work.

project with respect to the themes of this thesis, I first contrast Davidson's project with Lewis's. Whilst there are many commonalities, their respective views of the status of metaphysics are very different. I argue that the reason for this is that Davidson's revisionary views about the determinacy of interpretation invite a certain kind of structuralism about the content of metaphysical theories. Having established a robust role for truth in metaphysics, and a 'language-first' approach to metaphysical inquiry, Davidson's project demonstrates how the status of this metaphysical project can be transformed depending on claims made about other semantic concepts.

II. Truth, Meaning, and Interpretation

i. The Theory of Truth and the Theory of Meaning

In this section I first need to identify the scope of and constraints on Davidson's theory of meaning. Davidson holds that the primary aim of the philosophy of language is to understand natural language. (Davidson 1973a, 71) Davidson wants to offer a descriptive semantics for natural language. The aim of the theory of meaning is to be able to pair speaker's utterances with what they mean, in such a way that the meaning of complex expressions are composed out of the meaning of the simpler ones. The project is descriptive in that Davidson does not think that: "the task for a theory of meaning as I conceive it is not to change, improve, or reform a language, but to describe and understand it." (Davidson 1967, 29)

Davidson takes it as a condition on the acceptability of any candidate theory of meaning that it should be able to naturalistically explain how language is learnable. That is: how language speakers can understand complex statements given the apparent poverty of the empirical data to which they are exposed. Whilst Davidson thinks we should construct the theory of meaning in a way that starts from these reflections about learnability, Davidson does not think it

is appropriate to allow purely a priori considerations dictate our theories of language learning. (Davidson 1965, 3) How language is actually acquired Davidson rightly considers a matter for empirical study. However, he does think that philosophers are entitled to consider the conceptual matter of what gets to count as 'knowing a language'. (Davidson 1965, 7-8) Thus his arguments are intended to provide a framework for describing what language must be like if it is to be learnable, whilst leaving the actual process open for discovery by empirical linguists. In separating the empirical question from the conceptual one he is effectively demarcating what he sees as the limits of the empirical. (Davidson 1965, 3) A presupposition of this type of empirically orientated project is that Davidson does not think it is permissible to assume any sort of semantic knowledge of the sort we are trying to explain.

The primary consequence of this is that the semantic theory must be compositional. Language speakers are in principle able to understand indefinitely many sentences, but they only have finite cognitive capabilities and experience. Davidson holds that the best explanation for this is that language must have a recursive structure. So, the theory of meaning must account for the meaning of every sentence of the language by analysing it as composed "in truth relevant ways, of elements drawn from a finite stock". (Davidson 1970, 56) Davidson thinks that the best way to account for compositionality is to give an analysis of language structure, and that the best way to analyse structure is by offering a theory of truth:

"I suggest that a theory of truth for a language does, in a minimal but important respect, do what we want, that is, give the meanings of all independently meaningful expression on the basis of an analysis of their structure." (Davidson 1970, 55)

For Davidson, giving an account of compositionality does not then mean giving an account of the individual *meanings* of all expressions of the language. He claims that we should not think that individual words have meaning in any sense that transcends the fact that they have a systematic effect on the meanings of the sentences in which they occur. (Davidson 1967, 18) What Davidson wants to avoid is the view that an adequate theory of meaning must find entities to serve as the meanings of all the expressions in a language.

(Davidson 1970, 61) There are two motivations for this. First, Davidson thinks that any such entities are explanatorily superfluous, in that giving adequate rules for combination provides a sufficient basis for compositionality. He claims that this is done by offering a definition of truth in a language: “A theory of meaning for a language L shows ‘how the meanings of the sentences depend on the meanings of words’ if it contains a (recursive) definition of truth-in-L.” (Davidson 1967, 23) The second motivation is more methodological. In the spirit of an empirical semantics project, Davidson wants to situate the theory of meaning as closely aligned with an epistemology of linguistic understanding, and he doesn’t think that making sense of how a community of language users interpret and understand one another requires recourse to an ontology of intensional entities.

Having outlined Davidson’s aims and constraints on the theory of meaning, I now want to outline his positive proposal. The task for Davidson is to provide a semantic theory which is compositional and does not require antecedent understanding of semantic notions. So how does the Davidsonian project proceed? Davidson wants to take a semantic concept of truth and use it to “throw light on meaning, interpretation, and translation.” (Davidson 1977a, 204) So what kind of theory of truth does Davidson have in mind? He tells us:

“By a theory of truth, I mean a theory that satisfies something like Tarski’s Convention T: it is a theory that by recursively characterising a truth predicate (say ‘is true in L’) entails, for each sentence of L, a metalinguistic sentence got from the form ‘s is true in L if and only if p’ when ‘s’ is replaced by a canonical description of a sentence of L and ‘p’ by a sentence of the metalanguage that gives the truth conditions of the described sentence. The theory must be relativized to a time and a speaker (at least) to handle indexical expressions. Nevertheless I shall call such theories *absolute* to distinguish them from theories that (also) relativize truth to an interpretation, a model, a possible world, or a domain. In a theory of the sort I am describing, the truth predicate is not defined, but must be considered a primitive expression.”

(Davidson 1977b, 215-216)

A couple of comments on the notion of truth at work here is required. First, Davidson says that his theory is concerned with an absolute notion of truth, and indeed that this is the notion with which philosophers of language should be concerned. (Davidson 1973a, 69) He claims that those philosophers of language who aim to characterise a relativised concept of truth (e.g. T in a model, possible world, interpretation, etc.) are engaged in a different project. (Davidson 1973a, 68) The reason for this is that model theoretic accounts need to be able to take the intended model as a primitive of the theory, which imports the kind of semantic machinery to which Davidson doesn't think we should help ourselves. By contrast, Davidson's preferred satisfaction theoretic account takes truth to be satisfaction in all sequences, and so does not need to relativise truth to an intended model. Saying this, however, he wants to allow that Convention T may be partly relativised if it is to be applicable to natural language. (Davidson 1970, 58) Truth in a given language is "a relation between sentences, speakers, and dates." (Davidson 1969, 43) So whilst he wants an absolute notion of truth, which is characterised with only resources from the sentences of the language itself, he allows that, when applied to natural language, truth must be relativised to a speaker and a time. He maintains, however, that this notion of relativised truth differs from truth in a model, as the verification instances of the T-sentences "remains respectably empirical". (Davidson 1973a, 75)

The second significant thing of note about this semantic conception of truth, is that a semantic theory of truth is not a metaphysical theory of truth. Unlike a correspondence theory, which defines truth as a robust relation between representational stuff on one side and worldly stuff on the other, the semantic theory of truth does not tell us *what truth is*. Rather, it characterises a predicate which applies to a sentence in a language (true-in-L) just in case the truth conditions of a sentence in L depend upon the sentence's structure and the properties of its component parts. That is: a semantic theory of truth shows us how to generate for any sentence of a language, a T-sentence:

(T) 'S' is true iff P

So how do we get a theory of meaning out of this? Davidson thinks that understanding a language consists in knowing what it is for a sentence to be

true. (Davidson 1967, 24) Whilst truth conditions are not the same things as *meanings*, Davidson thinks that by giving the truth conditions of a sentence we give its meaning. (Davidson 1970, 56) The type of truth conditions that Davidson is after are not theoretically complex. He holds that the T-sentence must state “the truth conditions of a sentence using resources no richer than... those of the sentence itself.” (Davidson 1973b, 132) Thus, unless the original sentence mentions possible worlds, intensional entities, properties, or propositions, the statement of its truth conditions should not. Given that we antecedently understand the concept of truth, Davidson’s proposal is that we use this understanding to give the meaning theory. So, instead of something of the form:

(M) ‘S’ means that P

We rather have:

(T) ‘S’ is true iff P

Where ‘S’ is a sentence of the object language, and P is a sentence of the metalanguage. So, to substitute in a sentence of English:

(T1) ‘Snow is white’ is true iff snow is white

To see how this adequately gives the meaning of ‘S’ it helps to consider a heterophonic case, where the object language and the metalanguage are different. So:

(T2) ‘Schnee ist weiss’ is true iff snow is white

Two comments on this proposal are required. First, I claimed above that Davidson takes the primary constraint on a theory of meaning to be compositionality. We therefore require a theory where the meaning of whole sentences is dependent on the parts that make it up. Davidson makes good on this constraint because the Tarskian conception of truth employs a notion called *satisfaction*. Satisfaction is a recursive concept which stands to open sentences as truth stands to closed sentences, in that the satisfaction conditions of complex sentences depend on the satisfaction conditions on simpler sentences.

Second, Davidson thinks that it is important that we have some way of relating expressions of the language to things in the world in a way that “yields acceptable results with respect to the truth conditions of sentences.”

(Davidson 1977b, 224) He thinks that the relations between names and predicates and objects are given content indirectly when the T-sentences are given. (Davidson 1977b, 223) Thus whilst our semantic theory may use terms like ‘refers’, we should not take this mere theoretical posit to be doing any important empirical explanatory work. As long as we have some way of relating language to the world (and he thinks this will be in terms of the T-sentences) we have retrieved everything that we want from the notion of reference without introducing some kind of external semantically special relation. One of his motivations for this is that he thinks it is required for an empirically respectable theory. He thinks that we should account for the meanings of expressions not by trying to find referents for the individual words, but rather by considering the role played by the expressions in the whole language. (Davidson 1977b, 225)

The test of adequacy for the theory of meaning is then that it must be able to provide a T-sentence for every sentence of the language. (1969:46, 1973b:132) That is: for each sentence of the language we must generate a sentence which offers a statement of the condition under which the sentence is true. Davidson takes it that empirical verification of the meaning theory is a matter of demonstrating that the T-sentences are true. (1973b:133) This is therefore a holistic conception of linguistic understanding: the unit of empirical significance is the whole meaning theory.

ii. Radical Interpretation

In the last section I outlined the structure for Davidson’s theory of meaning. A requirement of the theory of meaning for Davidson isn’t merely that we can discern which expressions are *meaningful*, but that we can, for any arbitrary sentence of the language, determine what that sentence means. Davidson’s project of radical interpretation is an attempt to offer an account of this sort of

linguistic meaning in a language. In this section I outline Davidson's project of radical interpretation, and highlight two significant consequences of the project.

The challenge of radical interpretation is that we cannot determine the meaning of speaker utterances if we do not know what speakers believe, but in order to identify belief contents we need to presuppose a theory of meaning. The radical interpreter addresses this challenge by providing a theory of meaning and a theory of belief at the same time. A methodological assumption on the behalf of the interpreter is the principle of charity. Davidson claims that speakers hold "a sentence to be true because of what the sentence (in his language) means, and because of what he believes." (Davidson 1973b, 134) Charity can roughly be thought of a commitment to two further principles: coherence and correspondence. Coherence holds that when we are interpreting speakers of an alien language we should take them to reason in accordance with classical logical laws. Correspondence says that the interpreter should take there to be a relationship between the beliefs and utterances of speakers of a language and their environment. What these two principles amount to is that attributions of belief and assignments of meaning must be both consistent with one another and with the speaker's behaviour, and they must also be consistent with the evidence afforded by our knowledge of the speaker's environment. Davidson claims that if the interpreter cannot find a way of interpreting utterances in such a way as to reveal "a set of beliefs largely consistent and true by our own standards, then they have no reason to count that creature as rational, as having beliefs, or as saying anything." (Davidson 1973b, 137) An empirically adequate theory of meaning must generate a T-sentence for every sentence of the object language, whilst respecting the principle of charity. Radical interpretation then proceeds by holding belief fixed to work out meaning. (Davidson 1973b, 137) A toy example of the how radical interpretation might proceed is as follows. First the interpreter begins with observations such as:

- (1) x belongs to a community of speakers of a common language, K , and she holds "Gavagai!" true on Saturday at noon, and there is a rabbit visible to x on Saturday at noon,

By observing *K*-speakers' verbal behaviour, the interpreter gathers more evidence and infers that

(2) If *x* is a *K*-speaker, then *x* holds "Gavagai" true at *t* if and only if there is a rabbit visible to *x* at *t*.

Having gathered an adequate sample of instances of *K*-speakers holding "Gavagai" true when and only when rabbits cross their paths, the interpreter takes (2) to be confirmed. In turn, then, the interpreter takes (2) as evidence that (partly) confirms the following T-sentence of a Tarski-style truth theory for *K*:

(3) "Gavagai" is true when spoken by *x* at *t* if and only if there is a rabbit visible to *x* at *t*.

The interpreter continues adjusting this biconditional in light of new data, aiming to maximise agreement, whilst constrained by the formal requirements of the theory of meaning. Davidson takes this maximisation of agreement as a legitimate method as he thinks that we cannot attribute mass error or else we risk "depriving the subject of his subject matter." (Davidson 1977a, 200) Whilst the interpreter cannot attribute mass error, they will need to make allowances for explicable error or else run the risk of interpreting a speaker as only having true beliefs. Therefore, adjusting (3) consists in part in admitting that *x* may have some false beliefs, where the explanation for this may invoke environmental factors.

This invites a comment on two important things to note about this project and its set up. First, because of the methodological presuppositions for radical interpretation, we can, if we like, talk about the content of belief or sentences as corresponding to the way the world is. Built into the principle of charity was the claim that there is a relationship between the content of beliefs and utterances and the speakers' environment. Thus, in negotiating an interpretation to make allowances for explicable error, the interpreter will bring in environmental factors as explanans. This is absolutely not to say that Davidson offers a correspondence theory of truth. Rather, the position is that when we use the theory of truth to give the theory of meaning (in a Davidsonian style) and we are engaged in a radical interpretation project, a

result of these commitments is that there is a correspondence between the content of our utterances and the worldly things they are about. The motivation here is not hefty metaphysical considerations to do with the nature of truth, but rather naturalistically respectable constraints on the theory of interpretation.

A second point to note, that will become important in the next section, is that Davidson does not think it necessary to construct a theory of meaning which totally eliminates indeterminacy. He thinks that there can be a serviceable theory which does not uniquely pick out a best interpretation of speakers utterances. (Davidson 1974, 284) However, although he does not think we will always be able to identify a unique interpretation, he does not seem worried that this will lead to unacceptable indeterminacy. (Davidson 1977b, 224) I address whether Davidson's optimism about the acceptability of indeterminacy is well placed below.

III. Davidsonian Metaontology

iii. Reference and Ontology

In this section I offer a reconstruction of Davidson's approach to metaphysics. Davidson is explicit in several places about his view on the proper method of metaphysics. Whilst Davidson's views on truth, meaning, and first-order metaphysics (e.g. the metaphysics of action) have generated a large amount of secondary literature, there is little discussion of his metaontological views. In addition to helping forward the overall goal of this chapter, then, this section helps to fill a lacuna in the commentary on Davidson's work.

Davidson's views on the nature and method of metaphysical inquiry are inseparable from his views of truth and language. The background motivation for this view is that "successful communication proves the existence of a shared, and largely true, view of the world." (Davidson 1977a, 201) The central

notion at the work in Davidson's view of metaphysics is that of truth. He tells us that to say of a statement that it is true, is to say that it is true to the facts, where this is to be explained by means of an "appeal to a relation between language and the world". (Davidson 1969, 37-38) He claims that:

"Statements are true or false because of the words used in making them, and it is words that have interesting, detailed, conventional connections with the world. Any serious theory of truth must therefore deal with these connections, and it is here if anywhere that the notion of correspondence can find some purchase." (Davidson 1969, 43)

This talk of correspondence is not to be explicated in terms of reference. Davidson explicitly eschews talk of reference, but does not think that his rejection of reference is to dispense with ontology: "doing without reference is not at all to embrace a policy of doing without semantics or ontology." (Davidson 1977b, 223) Thus although he thinks that we do not need a notion of reference, there is room for a notion of correspondence, and he thinks that Tarski's semantic conception of truth can serve this purpose "because of the part played by the concept of satisfaction." (Davidson 1969, 48) He holds that the role played by 'corresponds to' and 'satisfies', although not identical, serves his metaphysical purposes as they both express a relation between language and the world, and both are equivalent to 's is true' when s is a (closed) sentence. (Davidson 1969, 48) Davidson claims that "in explaining truth in terms of satisfaction, all the conceptual resources of the language in relation to its ontology are brought to bear." (Davidson 1969, 49)

The positive proposal for how we should prosecute answering metaphysical questions is that we should proceed by a systematic study of language.

Davidson tells us that:

"One way of pursuing metaphysics is therefore to study the general structure of our language." (Davidson 1977a, 199)

The proposal is then that we should attend to the logical structure of the sentences *about* entities rather than attending to the entities themselves. How does this work? First, Davidson tells us that if we want to derive metaphysical conclusions from the theory of truth, our approach to language must be

holistic. (Davidson 1977a, 203) We must attend to what is required for a sentence of the language to be true but we need to do this for the whole language, not in a piecemeal way which proceeds sentence by sentence. Davidson's strong metaphysical claim is that "if the truth conditions of sentences are placed in the context of a comprehensive theory, the linguistic structure that emerges will reflect large features of reality." (Davidson 1977a, 201) We arrive at specific ontological commitments when the structure exposed is quantificational. That is: when the truth conditional analysis requires quantifiers and domains of quantification which include entities. If the sentences are true, then we are ontologically committed to the entities in question. (Davidson 1977a, 210)

This bears an interesting similarity to Sider's view of how ontological inquiry should proceed. There are of course differences, which I want to quickly highlight. Although they both think that a thoroughgoing analysis of the structure of language is the best method for deriving our metaphysical commitments, Sider is interested in the *fundamental* commitments. Sider is in a position to make this distinction between the superficial commitments and the fundamental ones as he thinks that when doing metaphysics we should not be analysing ordinary language, but rather pursuing inquiry in a precise language of metaphysics. Here, then, is the most significant difference between Sider and Davidson's proposals. Davidson is urging that the language that we need to analyse is the natural language of communities of language users. As we saw above, how he motivates the notion of correspondence is by considerations to do with the interpretation of the utterances of linguistic communities. In one sense then, Davidson is actually making a stronger claim than Sider about the relationship between the analysis of language and metaphysics: it is the language that we all actually speak from which we derive our metaphysical commitments, whereas Sider holds that if we are to work out our metaphysical commitments from language, we need to move our focus from natural language to a metaphysical language.

Davidson does not think that his proposed method will settle all matters of metaphysical interest. He claims that the:

“Merit of the method of truth is not that it settles such matters once and for all, or even it settles them without further metaphysical reflection. But the method does serve to sharpen our sense of viable alternatives, and gives a comprehensive idea of the consequences of a decision.” (Davidson 1977a, 214)

When discussing the abstract issues of the methodology of metaphysics Davidson does not give much indication of how the ‘further metaphysical reflection’ should be understood. However, if we consider his first-order metaphysical theorising, I suggest that the further metaphysical reflection does not consist in an independent *metaphysical* method. Instead, this will consist in reflecting on which statements we should accept as true (where this enterprise is continuous with the natural sciences), and on the proper method for analysing those statements.

iv. Semantic Indeterminacy and Ontology Relativity

Davidson wants his theory of meaning to be empirically respectable. The T-sentences are the empirical testable consequences of the theory of meaning, so if his theory of meaning is going to have empirical verification, it is purely in virtue of the verification of the T-sentences. Verification of the T-sentences (sentences such as “‘snow is white’ is true iff snow is white”) happens when there is mutual agreement of a linguistic community. Any foundational semantic machinery only gets verification inasmuch as the T-sentences do. For example, he tells us that, “meanings of words, reference, and satisfaction are posits we need to implement a theory of truth. They serve this purpose without needing independent confirmation or empirical basis.” (Davidson 1977b, 222)

An interpretation of language must be an interpretation of the whole language, and that interpretation gets verified only in so far as the T-sentences are. Davidson does not think that the constraints that we place on interpretation are sufficient for generating a best interpretation of the language. (Davidson 1974, 284) However, although he does not think we will always be able to hit

on a unique interpretation, he does not seem worried that this will lead to unacceptable indeterminacy:

“I do not for a moment imagine that such uniqueness would emerge. But I do think that reasonable empirical constraints on the interpretation of T-sentences (the condition under which we find them true), plus formal constraints, will leave enough invariant as between theories to allow us to say that a theory of truth captures the essential role of each sentence.” (Davidson 1977b, 224)

Whilst Davidson does not appear to be concerned about the lack of uniqueness of interpretation, I will show in this section that Davidson’s brand of indeterminacy, coupled with his view of how we generate ontological commitments, entail that he is committed to a type of Quinean ontological relativity. Davidson often acknowledges his philosophical debt to Quine. So, to understand Davidson’s brand of indeterminacy, I first must outline the Quinean brand of indeterminacy.

In Quine’s ‘Ontological Relativity’ a certain conception of meaning comes under attack. It is one according to which “Uncritical semantics is the myth of a museum in which the exhibits are meanings and the words are labels”. (Quine 1969, 27) Quine’s claim is not merely an epistemic thesis that we cannot know what our words refer to, but rather the stronger claim that there is no fact of the matter about what our words refer to. (Quine 1969, 47) Quine writes:

“When. . . we turn toward a naturalistic view of language and a behavioral view of meaning, what we give up is not just the museum of figure of speech. We give up an assurance of determinacy. Seen according to the museum myth, the words and sentences of a language have their determinate meanings... For naturalism the question whether two expressions are alike or unlike in meaning has no determinate answer, known or unknown, except insofar as the answer is settled in principle by people's speech dispositions, known, or unknown. If by these standards there are indeterminate cases, so much the worse for

the terminology of meaning and likeness of meaning.” (Quine 1969, 28-29)

To arrive at his conclusion Quine relies on the radical translator we meet in *Word and Object*. (1960) Quine invites us to consider how, starting from a position where we have only non-semantic evidence, we might come to know what the expressions of a given language mean. The radical translator’s job is to translate the language L of a distant linguistic community, whose history she knows nothing at all about, and whose linguistic practices and conventions may, for all she knows, be very different from hers, into her native language. The radical translator thus starts off from a point of complete ignorance about the community and its language. Since there are infinitely many expressions of L, she cannot specify them individually, but will have to provide (i) a translation manual which assigns meanings to each of the atomic parts of the language, together with (ii) a set of compositional semantic rules, that show how the meaning of complex expressions can be computed on the basis of their structure, and the meaning of their component parts.

The translator encounters a problem. That is: “a whole rabbit is present when and only when an undetached part of a rabbit is present; also when and only when a temporal stage of a rabbit is present.” (Quine 1960, 30) So, when the translator tries to translate ‘gavagai’, she has a number of eligible candidate translations available. There is no way to tell the candidate translations apart by ostension; there is no way to tell, purely on the basis of observing overt behaviour, how the alien language speaker is individuating. We may well prefer one interpretation over another, but this decision is purely at the level of pragmatics: “the maxim is [her] own imposition, toward settling what is objectively indeterminate. It is a very sensible imposition, and I would recommend no other. But I am making a philosophical point.” (Quine 1960, 34)

The moral to be learned from Quine’s radical translator is that we can’t talk about the *absolute* translation of an expression. Relative to one translation manual ‘gavagai’ means ‘rabbit’, relative to another ‘gavagai’ means ‘undetached rabbit parts’. There is no further fact which makes it the case that one or another of these translations is really getting it right. The only sort of facts that

could be the constitutive facts about meaning do not exist. The point Quine is trying to make is not about translation, but rather that viewed correctly our words are just as indeterminate as those of the alien linguistic community: “Radical translation begins at home”. (Quine 1969, 46) What legitimates the move from the indeterminacy in translation (from an alien language to our own) to indeterminacy about our own language itself is that the data we have available in the first-person case is just the same as the data we have in the third-person case. In both cases the ‘facts’ about meaning are exhausted by the behavioural facts. Quine claims that semantic reality disappears when one descends below the level of whole sentences. There is no fact of the matter as to which entity, out of the range of possible candidates, a particular term refers.

Quine claims we require a metalanguage in which to conduct our semantic inquiry. But of course, now situated in the background language, we can always probe more, and ask of the terms if *this* language: do the terms really refer to this or that. These questions, yet again only make sense to relative to a further background language. (Quine 1969, 49) We therefore regress, and we end up, according to Quine, back at direct ostension. As we found that ostension could not provide us with determinate answers in the case of the translator, we cannot expect any more satisfying answers now. Making an analogy with velocity, Quine makes the point that, like with questions pertaining to absolute velocity, questions about absolute reference do not make sense. Only relative to some specified system can we get a handle on an answer, or indeed a sensible question. (Quine 1969, 50) To ask such questions absolutely does not make sense.

However, it is crucial to understand that Quine is not saying that it is nonsense to talk sensibly and distinctively about rabbits and rabbit parts. What is fundamental to understanding Quine’s remarks is that such talk only makes sense relative to a frame of reference, or co-ordinate system. We can talk of meaning one thing rather than another once we have decided on a frame of reference. But this decision is pragmatic, and in some sense arbitrary. Quine tells us that this method is “laudable in practice and the best we can hope for”,

nonetheless, it does not provide us with a factual base for meaning. (Quine 1969, 33)

The lesson to be learned from the inscrutability of reference according to Quine was that it only made sense to talk of reference in terms of a system. However, what starts out as a concern with the (non)facticity of meaning develops into a concern about ontology. He writes:

“In the language of the theory there are predicates by which to distinguish portions of this universe from other portions, and these predicates differ from one another purely in the roles they play in the laws of the theory. Within this background theory we can show how some subordinate theory, whose universe is some portion of the background universe, can by a reinterpretation be reduced to another subordinate theory whose universe is some lesser portion. Such talk of subordinate theories and their ontologies is meaningful, but only relative to the background theory with its own primitively adopted and ultimately inscrutable ontology.” (Quine 1969, 51)

What should be seen as the metaphysical significance of Quine’s comments here is that theories can never be fully interpreted, except in the relative sense. (Quine 1969, 51) A theory, according to Quine, is a fully interpreted set of sentences. However, for the theory to be fully interpreted, the values of the bound variables must be settled. How do we settle this? If we accept what Quine has already told us, we cannot fix the objects of a theory absolutely. We do so only relatively. Whilst we can talk of ontology, we can only do so relative to some background theory, which can itself, in turn, be questioned and explained in terms of some other one. That is: “There is no absolute sense in speaking of an ontology of a theory.” (Quine 1969, 60)

Davidson explicitly agrees with Quine that this type of *semantic* indeterminacy occurs:

"Quine is right ... in holding that an important degree of indeterminacy will remain after all the evidence is in; a number of significantly different theories of truth will fit the evidence equally well." (Davidson 1970, 62)

However, there is a significant divergence between Davidson's and Quine's views on the consequences of this indeterminacy. Davidson says he is not worried about the indeterminacy; that it is not a type of indeterminacy that we should be worried about. He draws a comparison with measurements of temperature. We have three distinct ways of measuring temperature (and indeed there could be more). In the temperature case and in the case of interpretation, there are some empirically observable phenomena, and there are different theories that assign different values to the objects at issue. However, there is no difference in the empirical adequacy of those theories. What matters, Davidson claims, in the overall pattern of assignments, rather than individual cases.

If all we are concerned with is getting the right distribution of truth values for the sentences that make up a theory, then perhaps we can agree with Davidson that such indeterminacy is not problematic. However, if when doing metaphysics we care about more than mere verificational equivalence, then I think we have reason to be concerned. Recall, we incur ontological commitments when we engage in a systematic study of the language, and when this study reveals true quantified statements and a domain of quantification. The sort of subsentential indeterminacy that Davidson admits into his theory of interpretation means that we cannot talk absolutely or determinately about the values of the bound variables in the scope of the quantifiers.

If we favour Davidson's approach to metaphysics are we compelled to accept some kind of non-realist position about the things to which we are ontologically committed? Maybe we must just accept that there isn't a fact of the matter about whether we are committed to rabbits or undetached rabbit parts. This isn't devastating if you are happy with the thought that metaphysics should not be in the business of offering theories which tell us what the world is determinately like. A possible alternative which I think merits further investigation is whether we could think about Davidson as being committed to a realism about certain ontological *structures*. What matters is the overall pattern of assignments, so if you favour ontological realism perhaps this could be explicated as a sort of structural realism. Ontic structural realism is the view that the world has a certain structure, "in the sense of not supervening on the

intrinsic properties of a set of individuals”, and it is that structure to which we are ontologically committed.² (Ladyman & Ross 2007, 104) It would take me too far afield to fully develop this structuralist reading of Davidson, but it nonetheless gestures towards a potentially fruitful avenue of investigation for anyone who likes Davidson’s approach to metaphysics but wants to resist any non-realist implications of ontological relativity.³ What matters for my current purposes is that however the implications of ontological relativity are understood, Davidson’s metaphysics cannot provide the sort of fundamental and realist metaphysical theories at which the projects in the previous three chapters aim.

IV. Truth and Metaphysics

v. Comparisons with Lewis

The purpose of this section is to critically compare the semantic and metaphysical frameworks offered by Davidson and Lewis. I argue that both philosophers give a central and integrated role to truth in their semantic and metaphysical theorising. However, whilst their big picture commitments display similarities, there is much divergence in the detail. The purpose of this section is to help to understand how the space of possible responses to Davidson’s indeterminacy and consequent ontological relativity is constrained by his views on interpretation.

² Debates about ontic structural realism are almost exclusively situated in the philosophy of physics. The application of this approach to the kind of metaphysical project I am considering here would therefore require some substantial work.

³ In his (2008) paper, Ritchie develops a position he calls ‘semantic structuralism’ which draws on Davidson’s theory of meaning. Whilst Ritchie’s project is directed at a different set of concerns from those I consider here, this paper provides a good basis for thinking about structuralist readings of Davidson.

Davidson and Lewis share a commitment to a truth conditional theory of meaning. They agree that the best handle we have on 'meaning' is in terms of truth conditions; that we should give the meaning of expressions in the language by offering their truth conditions. But their projects come apart. Whereas Davidson thinks the truth conditions should take the form of a Tarski style T-sentence, Lewis is in the business of offering world invariant truth conditions. However, whilst there is significant divergence on the proper specification of the truth conditions (Tarski T-sentences versus world-invariant truth conditions), they agree that to give an account of meaning you must be dealing with truth conditions. In addition to truth conditions playing a substantial role in the theory of meaning, they agree that truth conditions also play a role in metaphysics. In both cases, the truth conditions specified by the semantic theory orientate the metaphysical inquiry: the metaphysics is in some sense constrained by the semantics.

Despite these similarities, Davidson and Lewis end up with markedly different conceptions of the status of metaphysical inquiry. They agree both that the analysis of language is a prerequisite to metaphysical inquiry and that having offered a satisfactory analysis of the true statements of the language there will be things of metaphysical interest left unsettled. However, how Davidson and Lewis think we settle those outstanding issues of metaphysical interest is a first place to note that they come apart. Lewis thinks that we need to offer an independent metaphysical methodology for building the metaphysical theories which act as the satisfiers of the analysis of the language. There is no suggestion in Davidson that he wants to employ such methods. The primary dissimilarity which I want to exercise in the section, however, is their respective views about the determinacy of content. Although they also have different conceptions of how metaphysics unfolds, the more significant difference in their metaphysical projects stems from their different views on interpretation.

Davidson and Lewis engage in a similar interpretationist project. To exercise the differences in their approaches, I want to look at the constraints they place on radical interpretation. I take it this highlights fairly fundamental differences in how they think it is permissible to generate an interpretation of a language. They both agree that some narrowing down of the space of possible

interpretations is necessary, and therefore introduce constraints on what can count as an acceptable interpretation. However, these broadly empirical or behaviouristic constraints still aren't sufficient to land on a best interpretation. There is still room for indeterminacy (as we saw in the discussion in the last section). Davidson is happy to allow this indeterminacy, whereas Lewis wants to introduce a further constraint to eliminate what he sees as problematic indeterminacy.

Lewis claims that "there are various ways for a system of belief to make a belief sentence true." (Lewis 1986b, 32) Lewis doesn't side with Davidson that the interpreter should interpret as many beliefs/utterances as possible as coming out true. Lewis, by contrast, thinks that interpretations should *rationalise* agents. We should interpret someone so as many of their beliefs as possible some out rational. He wants to rule out cases where a subject is surrounded with misleading evidence, where we should, he thinks, interpret them as having false beliefs rather than just being epistemically lucky. Lewis thinks that the content of belief is determined according to the principles of instrumental rationality. (Lewis 1986b, 36) This claim for Lewis, similarly to Davidson, isn't merely about weeding out unacceptable interpretations of a speaker. He takes the principles of instrumental rationality to be part of what it is to have a belief. He tells us: "We should take this principle of instrumental rationality to be neither descriptive nor normative but *constitutive* of belief. It enters into the implicit definition of what it is for someone to have a certain system of beliefs." (Lewis 1986b, 36)

So far I have claimed that Davidson and Lewis introduce constraints on what gets to count as an acceptable interpretation. In both cases, although the constraints themselves differ, they think that their constraints characterise what it is to have a belief. However, even with these constraints in place, they both also agree that the kind of evidence available to the interpreter will seriously underdetermine content. Lewis tells us that: "Constitutive principles of fit which impute a measure of instrumental rationality leave the content of belief radically underdetermined." (Lewis 1986b, 38) In the face of this indeterminacy, they once again diverge. As I discussed in the previous section, Davidson thinks that this degree of indeterminacy can be accepted. As long as

we are getting the right kind of pattern of truth across the language and delivering the T-sentences for each sentence of the object language, then we have done all we need to do. We can accept that content will be underdetermined.

Lewis, however, disagrees. Lewis introduces some metaphysical machinery to help deal with the problem. He tells us that some candidate interpretations are more natural than others, and are thus more eligible interpretations:

“These beliefs and desires are unreasonable; though if twisted desire is combined with correspondingly twisted belief, then it may be that the failing lies entirely outside the purview of the department of instrumental rationality. So I say that other departments of rationality also may have a constitutive role. What makes the perversely twisted assignment of content incorrect, however well it fits the subject’s behaviour, is exactly that it assigns ineligible, unreasonable content when a more eligible assignment would have fit the behaviour equally well.” (Lewis 1986b, 38-39)

It should be noted that to secure determinacy Lewis has not resorted to some like ‘the semantic intentions of speakers’, but rather claims that the world has a robust role in determining the content of beliefs. It is in part his commitment to reference which enables him to make this move. Lewis thinks that reference isn’t merely a matter of what we do in language or thought when we refer, but also consists in eligibility of the referent. (Lewis 1983b, 47) Lewis’s means of securing determinacy therefore requires commitment to a metaphysical thesis that there are natural properties and that these properties are more eligible to serve as the referents of our utterances. Having secured determinacy Lewis can then go on and talk sensibly about the content of metaphysical theories.

Davidson, however, in the absence of a method for determining the intended interpretation of a language or theory, can only talk about the verificational adequacy of metaphysical theories. In sum: Davidson’s project bears many similarities with Lewis’s project. A necessary prerequisite to metaphysics is the analysis of language. However, unlike Lewis, Davidson is not in a position to secure subsentential semantic determinacy, and this has a rather radical impact upon the status of the metaphysical enterprise.

V. Conclusion

In this chapter I have articulated Davidson's view of the role of truth in metaphysics, and the subsequent status of metaphysical inquiry. I argued that Davidson assigns an important role to truth in metaphysics. Davidson advocates for an approach to ontological inquiry according to which we derive ontological commitments when the truth conditional analysis of the true sentences of the language requires quantifiers and domains of quantification. The concept of truth that plays this robust role in metaphysics is one which is shaped with the demands of an empirical semantics project in mind. Importantly, the nature of Davidson's empirical semantics project is such that we cannot ensure subsentential semantic determinacy. Thus, we cannot *absolutely* talk about the ontological commitments of our best theory.

I take Davidson's project to illuminate how the status of metaphysical inquiry can be transformed depending on claims made about other semantic concepts. If you think both that a truth conditional analysis of language is a prerequisite for generating metaphysical commitments, and that the interpretation of that language is indeterminate, then these commitments together have a serious impact upon the type of metaphysical project in which you can engage. In particular, you cannot make claims about *the* ontological commitments of a theory. For approaches to metaphysics which take the study of language to be in some sense authoritative, Davidson's project is instructive in demonstrating how assumptions about the nature of truth and the theory of meaning can transform the limit or status of metaphysics. In particular, in this context, the claim that content is determinate is not metaphysically innocent.

Chapter Five: Bob Hale, Amie Thomasson, and Truth for Metaontological Deflationists

I. Introduction

This chapter investigates the prospects for a deflationary approach to the project of ontology. The approach here considered, which I'll call *metaontological deflationism*, promises significant epistemological and methodological virtues. On the one hand, the approach aims to demystify the method of metaphysics. Unlike many of the other approaches to ontological inquiry considered thus far, metaontological deflationism does not require the use of any distinctly metaphysical machinery to answer ontological questions. Further to this virtue, and indeed because of it, these methods often put us in a position to give definitive answers to ontological questions. Thus, if a viable approach, metaontological deflationism demonstrates how we can have a realist ontology with a transparent epistemology.

Truth plays a significant role in the metaontological deflationist framework: truth is that feature which the sentences from which we derive our ontological commitments must have. I argue that for the metaontological deflationist approach to be viable, for it to be able to deliver on all it promises, it must articulate an adequate conception of truth. I contend that as it stands the approach is lacking such an articulation. Whilst a significant theoretical role is assigned to the concept of truth, in order to sustain this role, the metaontological deflationist needs to provide some kind of story about how it is that truth can perform the function required of it. My discussion focuses on two approaches to metaontological deflationism: the neo-Fregeanism of Bob Hale and the neo-Carnapian project of Amie Thomasson.

The chapter proceeds as follows. First, I consider Hale's neo-Fregeanism. (§II) I situate the neo-Fregean approach to the metaphysics of mathematics, and

then articulate Hale’s development of this approach into a general approach to ontological inquiry. Second, I consider Thomasson’s neo-Carnapianism. (§III) I introduce her project as a direct response to ‘hard’ approaches to ontology, and then outline her proposed ‘easy’ method and its consequences. Finally, I turn to assess the role of truth in their projects. (§IV) I argue that both approaches to ontology are unified by their endorsement of three constraints on the metaontological deflationist approach: first, that truth must be capable of playing a robust theoretical role; second, that the ontological commitments delivered by such an approach be construed realistically; and third, that the ontological commitments should not be *metaphysically* explanatory. In trying to adhere to these constraints, I argue that truth has an even more important role to play in these approaches than initially appears to be the case: it is only when the metaontological deflationist has spelled out an adequate conception of truth that these commitments can be maintained together. I close by recommending a way of thinking about truth for the metaontological deflationist that satisfies their requirements.

II. Hale

i. Introducing neo-Fregeanism

Hale and Wright (2009) set up neo-Fregeanism as a response to an epistemological challenge in the philosophy of mathematics. As they present it, the problem is this: how can our mathematical knowledge fit into a naturalistic conception of knowledge acquisition? (Hale & Wright 2009, 178) Given the abstract nature of mathematical entities, we need some kind of story about how we can come to know truths pertaining to mathematical entities, even though we cannot enter into direct causal relationships with them.¹ The move

¹ One needs to be sensitive with this claim. The neo-Fregean does not mean to endorse a radically naturalistic epistemology of mathematics. Given that they hold that mathematics is a

proposed by the neo-Fregean in the face of such a challenge is to claim that mathematical knowledge is essentially logical knowledge. Hale and Wright argue that we can achieve this kind of justification of mathematical knowledge by attending to the semantics of such mathematical statements. They claim:

“It is not that, *before* we can understand how knowledge is possible of statements referring to or quantifying over the abstract objects of mathematics, we need to understand how such objects can be given to us as objects of acquaintance or how some other belief-forming mechanisms might be sensitive to them and their characteristics. Rather we need to tackle directly the question how propositional thought about such objects is possible and how it can be knowable. And this must be answered by reference to an account of how meaning is conferred upon ordinary statements that concern such objects, an account which at the same time must be fashioned to cast light on how the satisfaction of truth conditions it associates with them is something that is accessible, in standard cases, to human cognitive powers.” (Hale & Wright 2009, 178-179)

Their claim is that we have made a mistake if we request prior justification of mathematical knowledge. Rather, it is via attending to the semantics for statements which contain mathematical vocabulary that an adequate epistemology can be offered. This claim is found explicitly in Frege, and it finds its home in the context principle. It is by attending to the semantic contribution made by number terms to the sense of propositions that we can gain an understanding of how such objects can be given to us as objects of knowledge. As Frege says:

“How, then, are numbers to be given to us, if we cannot have any ideas or intuition of them? Since it is only in the context of a proposition that words have meaning, our problem becomes this: To

priori, a naturalism which only tolerated a posteriori knowledge would preclude the possibility of knowledge of mathematics. They therefore accept that there can be a priori warrant for the claims of mathematics, but what must be naturalised is how this mathematical knowledge can be applied to gain knowledge of the empirical world. Thus, the challenge for the neo-Fregean is to offer a theory which accounts for the fact that mathematics in both a priori, *and* can be used to produce knowledge about the natural world.

define the sense of a proposition in which a number word occurs.”

(Frege 1980, §62)

Hale and Wright claim that “abstraction principles are the key device in the epistemological project so conceived.” (Hale & Wright 2009, 179) An abstraction principle is a stipulative implicit definition of a new operator, the Σ -operator, where Σ -terms are new terms introduced into the language.² (Hale & Wright 2009, 179) An abstraction principle takes the following form:

$$(\forall a) (\forall b) (\Sigma(a) = \Sigma(b) \leftrightarrow E(a,b))$$

The abstraction principle is intended to show that the truth-conditions of Σ -identities are coincident with the truth conditions for a kind of statement that we have prior understanding of and already know how to use. (Hale & Wright 2009, 179) The thought is that we can use this prior ability to gain knowledge of the identities of the referents of the Σ -terms, where the existence of those entities is guaranteed by the truth of the abstraction principle. Hale and Wright claim that it is,

“Permissible to fix the truth-conditions of one kind of statement as coinciding with those of another – ‘kind’ here referring to something like logical form – in such a way that the overt *existential* implications of the former exceed those of the latter, although the *epistemological* status of the latter, as conceived in advance, is inherited by the former.” (Hale & Wright 2009, 180-181)

Whilst some mathematical terms can be explicitly defined, the cases in which the abstraction principle needs to do work are where this cannot be done. The salient principle is an implicit definition of the cardinality operator, Hume’s Principle:

$$\text{HP: } (\forall f) (\forall g) (N(f) = N(g) \leftrightarrow E(f,g))$$

² One initial potential concern that Hale and Wright consider pertains to the general status of implicit definition. Hale and Wright take it that such concerns are concerns for implicit definition in general, rather than anything to do with abstractions in particular. They claim, then, that they will take “implicit definition [to be a] *default* legitimate practice.” (Hale & Wright 2009, 192) For the purposes of this chapter I follow them in this assumption.

Or: The number of *F*s is identical to the number of *G*s if and only if *F* and *G* are equinumerous.

Hume's Principle fixes the sense of the cardinality operator by stating the conditions under which statements involving the related singular terms are true. So, how does this answer the epistemological worry with which we started? We know that every instance of the biconditional is true, and, when we know that the right-hand side of the biconditional is true we therefore know that the left-hand side is true, because of the logic of biconditionals. So, the knowledge that Hume's principle is true, together with the knowledge that the right-hand side is true, provides the requisite justification for our knowledge that the left-hand is true.

Hale and Wright consider a challenge according to which the success of the implicit definition cannot be secured without some kind of metaphysical assistance. We are supposed to imagine a possible scenario where a particular abstraction will fail because the world just isn't the way it needs to be for the abstraction to be true. Thus, the challenge goes, it is not sufficient to simply stipulate the implicit definition, the world must be a certain way. (Hale & Wright 2009, 193) They claim that such a view is mistaken: a *good* abstraction alone is sufficient to undermine this "epistemically metaphysical possibility".³ (Hale & Wright 2009, 193) This is because the truth of the right-hand side of an instance of a good abstraction is *conceptually* sufficient for the truth of the left-hand side. Hale and Wright, therefore, claim that there is no metaphysical assistance required. (Hale & Wright 2009, 193)

The abstraction principle alone is sufficient for providing truth conditions for the introduced term. Thus, they claim, the implicit definition is sufficient for giving the new term a *sense*. (Hale and Wright 2009, 196) We may now ask: what secures the reference of the new operator? Hale and Wright claim that the "the abstraction operator refers (to a function) only if the singular terms it enables us to form refer (to objects)." (Hale & Wright 2009, 201) So what guarantees that that the singular terms refer? Hale and Wright claim that all

³ The question of what counts as a good abstraction is a serious and fraught question, and one with which I shall not engage here.

that is required for a singular term to refer is that it appears in a statement where it occupies a “*reference-demanding* position.” (Hale & Wright 2009, 202) I take ‘reference-demanding’ to require two conditions are met. First, that the singular term appears in a statement where it functions as a singular term, and second that that statement in question is true. The latter requirement is crucial. This is what secures the independence of truth conditions and existence.

The question then becomes: how can we guarantee that objects of the relevant kind exist?⁴ All instances of Hume’s Principle are true, and the left-hand side of the biconditional will be true in all those cases where the right-hand side is true. When the right-hand side of the biconditional is true, referents are secured as the interpreted second order logic that Hale and Wright endorse comes with a rich ontology of functions. Thus, the metaphysical question is answered as the truth conditions for the identity statements containing the singular terms formed from the cardinality operator (left-hand side) are coincident with those on the (right-hand side) which comes prepared with the needed ontological commitments. We might wonder whether this is sufficient. However, Hale and Wright characterise any further demand for reason to believe in the truth of the abstraction principle to be sceptical in spirit. (Hale & Wright 2009, 204) Whilst I suspect that this may not be sufficient to dispel any such sceptical worries in the philosophy of mathematics, I will now leave these matters and turn to examine how such an approach to the metaphysics of mathematical entities can be generalised into a unified approach to ontological inquiry.

ii. Neo-Fregeanism Generalised

In this section I consider how the neo-Fregean approach outlined above generalises into an all-purpose approach to ontology. Hale claims:

⁴ There are going to be multiple ways to respond to this question (i.e. multiple interpretations of second order logic) but all that I consider here is the proposal in Hale and Wright (2009).

“... defend a broadly Fregean version of the view that questions about what kinds of things there are are inseparable from, and in one way posterior to, questions about the logical analysis of language.” (Hale 2013, 8-9)

According to Hale the central question of ontology is what kinds of things are there? (Hale 2013, 9) To ask what kind of things there are requires, according to Hale, that we have a prior categorisation of expressions of the language. (Hale 2013, 10) The motivation for thinking of ontology in this way stems from a desire to have a neutral and general conception of what objects and properties are from which we can begin inquiry. There is a particular methodological challenge here. Hale is interested in how we are to understand the general category of *object* and *property*. He suggests that if we are to get a handle on these general concepts in a non-question begging way, we need to consult the logical structure of language. (Hale 2013, 18-19)

Imagine, for example, that we define ‘object’ as that which has only material parts or that which is apt to enter causal relations. We can see immediately that this would bias any subsequent metaphysical inquiry as it would force us to judge that there can be no objects which are causally isolated from us (like, for example, numbers). Hale, therefore, takes the neo-Fregean approach to this question of categorisation to be the only viable one, which won’t prejudge the outcome of ontological inquiry. This is clearly a very strong commitment. Given that Hale is claiming that this is the *only* acceptable approach to ontological inquiry, then, if it turns out that this neo-Fregean project cannot be prosecuted, it seems we must conclude that ontological inquiry itself cannot be sensibly conducted.

Hale’s recommended approach transforms the structure of how such inquiry precedes: “under the Fregean approach, questions about the existence of entities of this or that kind are transformed into questions about *truth* and *logical form*.” (Hale 2013, 14) So how do we decide what objects are? Hale tells us that “objects are the (typically) non-linguistic correlates of the devices of *singular* reference.” (Hale 2013, 11) Extending this thought to the case of properties

and relations, properties and relations are what (one-or more-place) predicates stand for.^{5,6} (Hale 2013, 12)

All we are committed to thus far is that we are going to derive the categories of ontological interest from the logical structure of language. There is going to be an outstanding question of what the particular objects and properties are to which we are going to be committed. That is: which given species of entities fall under those categories. Hale takes it that a Fregean answer to the question of categorisation suggests a Fregean answer to the latter question of what kinds of things exist. (Hale 2013, 13-14) He claims:

“If entities belonging to a certain ontological category just are what expressions of the corresponding logical category stand for, then we can argue for the existence of entities of that kind by arguing that there are true statements involving expressions from that category.” (Hale 2013, 14)

The thought here is that it is sufficient for a singular term to refer to something, if it functions as a singular term in a token statement, and that that statement is true.⁷ (Hale & Wright 2009, 185) Such statements could not be true, so the thought goes, unless their ingredient expressions successfully discharge their semantic function. Discharging their semantic function consists in having reference, and for a singular term to have reference it must refer to an object. (Hale 2013, 36)

You might worry about the potential expressive limitations of natural language for an approach such as this. If we want to say that objects are what the singular terms stand for, then we might be inclined to think that all objects

⁵ Hale offers a reformulation of this in light of the concept horse objection. Thus: “*Objects* are what are or could be *primary* semantic values of singular terms. *Properties* are what are or could be *primary* semantic values of predicates and so on.” (Hale 2013, 32)

⁶ It is worth noting that Hale claims to be happy with the thought that we may be able to carve out other conceptions of objects or properties from this general one. For example, we may restrict ourselves to talking just about the properties with which fundamental physics is concerned. Importantly though, this must be seen as a restriction of the general category and not the identification of a more fundamental kind of category.

⁷ There is going to be some kind of restriction here on logical form of such expressions, for example, cases of negative existentials.

have names. It seems plausible that not all objects have actual names. Hale manoeuvres around this concern by offering a modal account of language, which “transcends the contingent limitations of actual languages by drawing upon their possible extensions.” (Hale 2013, 20) Given this, the final position is:

“There exist objects of the specified kind if and only if (i) there are or could be singular terms which would, if non-empty, refer to objects of that kind and (ii) if there were such terms, there would be true atomic statements containing them.” (Hale 2013, 37)

You might also worry that that this way of thinking about the ontological project necessitates some kind of anti-realism. That because we require the prior logical categorisation of the language before we can say anything about the ontological categories, that this makes the categorisation in some sense language dependent and so, in turn, mind-dependent. Hale thinks this would be a mistake. Whilst our identification of certain expressions as being of a particular type (e.g. singular term, n-place predicate, etc.) is derivative upon the logical analysis of language, the truth-values of statements containing such expressions need not be, and mostly won't be, language dependent. (Hale 2013, 19) It is because we require the *truth* of the relevant statements that the entities to which we are committed need not be language dependent.⁸

Finally, you might worry that whilst the ontology to which we are committed is not mind-dependent in any sense, there is something lightweight about these commitments. Hale himself claims that such a view of ontological inquiry is metaphysically deflationary and lightweight. (Hale 2013, 41) We need to be careful with this claim. I suggest that we must separate two kinds of ways that we could say that ontology is lightweight: we might think that the things to which we are ontologically committed are themselves lightweight, or we might think that the conception of the project of ontology is lightweight. Part of what

⁸ One point worth mentioning here is that although the project does not commit us to anti-realism, it is compatible with it. So, for example following Dummett, we could think that the correct truth-predicate for the language is a non-bivalent one, and that this is sufficient to make a view anti-realist. If we endorse this epistemically constrained notion of truth, merely saying that the sentences of the language are true won't get you to realism. More on this later.

the neo-Fregean programme was supposed to embrace was mathematical Platonism, which is certainly not a ‘lightweight’ position (in the former sense). Thus, I take it that Hale takes a deflationary attitude towards the *project* of ontology, in that there is no distinctly metaphysical, autonomous methodology for ontology. Indeed, Hale’s claim is stronger than this: there could not be a viable method for ontology other than the deflationary neo-Fregean account he proposes.

III. Thomasson

iii. Situating ‘Easy Ontology’

Amie Thomasson advocates a view of ontological inquiry according to which ontological questions are not deep questions. She claims:

“I call this the ‘easy’ approach to existence questions, since it entails that those existence questions that are meaningful are not deep and difficult subjects for metaphysical dispute, but rather questions to be resolved straightforwardly by employing our conceptual competence, often combining this with empirical investigations.” (Thomasson 2015, 20)

Consider a particular disputed class of entities. We start by taking an “uncontroversial truth... and reasoning by what seem like trivial steps... to reach ontological conclusions.” (Thomasson 2015, 21) The uncontroversial truth might be: (1) there are two mugs on the table. From (1) Thomasson thinks we can derive (2) analytically: (2) the number of mugs on the table is two. Thus, we are going to be committed to numbers. The thought driving this commitment is that we cannot be committed to the literal content of (1) *and* deny that (2) is entailed by it. I will further elucidate this conception of ontological inquiry in the next section, but first I want to situate Thomasson’s project relative to both neo-Fregeanism and what might be seen as a more

dominant methodology of contemporary analytic metaphysics, and to explain the Carnapian inheritance of her approach.

Thomasson says that neo-Fregeanism *as it stands* is a manifestation of easy ontology.⁹ She also presents Stephen Schiffer on 'pleonastic entities' as another already-existing case of easy ontology. (Thomasson 2015, 132-134) An approach to ontology is easy if it has these two features: (1) all well-formed existence questions may be answered by conceptual and/or empirical methods, and (2) some disputed existence questions may be answered by means of trivial inferences from uncontroversial premises. (Thomasson 2015, 128) Thomasson charts the differences between these views as follows: the neo-Fregean requires true identity statements, Schiffer requires uncontroversial true statements, Thomasson requires that we use our conceptual competence and empirical skills. (Thomasson 2015, 141) So each approach is a more generalised form of the last. Her claim is that both of "these approaches have seldom been seen as part of a unified approach to existence questions, nor have the relations among these views been made clear". (Thomasson 2015, 20) Trying to make good on both these limitations is her ambition for her proposed methodology.

Thomasson claims that in the latter half of the 20th Century there has been a proliferation of ontological questions and theories. She attributes this trend to the growing dominance of what she terms the neo-Quinean approach to metaphysics.¹⁰ This neo-Quinean approach takes it that "the enterprise of ontology is that of developing the best 'total theory', where that involves choosing the theory with the most theoretical virtues." (Thomasson 2015, 2) Thomasson claims that such a view of metaphysics has become so dominant as to seem "natural, even inevitable". (Thomasson 2015, 3) It is Thomasson's contention that such an approach to metaphysics is not inevitable, neither historically nor philosophically. She wants to return, therefore, to a pre-

⁹ It is worth noting here that Thomasson is talking about Hale and Wright's neo-Fregean programme in the philosophy of mathematics, not Hale's more general programme outlined above.

¹⁰ The 'Quine' in 'neo-Quinean' isn't supposed to commit her to a view about the correct exegesis of the historical Quine.

Quinean conception of metaphysics, which sees metaphysics as primarily engaged in conceptual work. (Thomasson 2015, 11)

Thomasson presents a pressing question for a metaphysician: given the success of natural science as a means of explaining the world, what are the proper tools and aims for metaphysics? She claims that the historical division of labour between philosophy and science is one according to which philosophy and science are viewed as being distinct and complimentary. Such a view takes it that philosophy is primarily engaged in conceptual work, whereas science employs empirical methods and concerns matters of fact. (Thomasson 2015, 7) Favourably quoting Wittgenstein, she suggests that “we must do away with all *explanation*, and description alone must take its place.” (Thomasson 2015, 8) The putative virtue of this move towards a metaphysics concerned with description rather than explanation is that it provides a clearer and more modest conception of the methods, roles, and epistemology of metaphysics. (Thomasson 2015, 11) Thomasson positions her preferred project as neo-Carnapian. One point worth noting, however, is that whilst Carnap is primarily concerned with more formal languages, Thomasson is interested in ordinary English.¹¹ (Thomasson 2015, 44) This focus on ordinary language means that conceptual analysis (i.e. working out what the rules that govern our terms actually are) plays an important role in Thomasson’s project. (Thomasson 2015, 44)

Given that Thomasson is making use of a Carnapian approach she assigns an important role to the distinction between internal and external questions. (Carnap 1950) She follows Price (2009) in claiming that the best way to interpret the Carnapian internal/external distinction is in term of the use-mention distinction. (Thomasson 2015, 36) So, when we ask internal questions we must use the relevant vocabulary, whereas when we ask external questions we merely mention such terms.

¹¹ It is worth noting that this move to ordinary language presents a particular kind of challenge. Ordinary language is an incredibly complex phenomena, and some have argued that this makes Thomasson’s project unworkable. I do not engage this challenge at present, but see Button (2016) for some interesting discussion.

On internal questions, Thomasson says that we can deliver simple answers to such questions. She claims:

“The very rules for introducing property language (combined with ‘customary deductive rules’) licence us to infer from an ordinary truth like ‘the house is red’ that ‘the house has the property of being red’ and so to provide an easy affirmative answer to the general existential question (asked internally) ‘Are there properties?’” (Thomasson 2015, 37)

She urges that we interpret the Carnapian claim as the claim that in order to use terms meaningfully there must be rules governing their use. (Thomasson 2015, 39) That is: to ask an existence question, the suspect term must be associated with some rules for use in a language. However, once we’ve got a context in which to ask and answer the existential (internal) question, the truth of the existential statement need not be thought of as truth-in-a-framework. (Thomasson 2015, 39) Once we adopt a language in which to ask existential questions, the answers that we derive, the existential statements that we hold to be true, are neither arbitrary nor language dependent. (Thomasson 2015, 43)

By contrast, the metaphysician who asks the external question attempts to cut those terms off from the internal rules which govern their use, and they therefore render the question meaningless. (Thomasson 2015, 40) But we can recover some sense for the external question if we understand them as questions which *mention* the suspect terms and ask whether we should adopt such terms with their associated rules for use. (Thomasson 2015, 40)

Thomasson claims that this is a charitable interpretation of hard ontologists like the neo-Quinean, as we do recover some reason to engage in those debates. The question of which linguistic framework to adopt is itself a practical question, but Thomasson claims that it may be influenced by theoretical considerations. (Thomasson 2015, 42) So, “the rules we adopt need not be arbitrary, given our purposes, since some rules may serve the purposes better than others.” (Thomasson 2015, 43) But, importantly, the debate about which linguistic framework best serves our theoretical or practical purposes should not be viewed as a debate whether or not there *really are* entities of a particular kind.

I take it that the kind of argumentative move that Thomasson is trying to avoid here is the following. Imagine that all parties agree that some discourse is useful, and further they agree that the true sentences of that discourse commit us to, for example, the existence of properties. What Thomasson wants to preclude is that there could be some further *metaphysical* question, which is distinct from whether the discourse is useful (external) or whether there are true statements containing the suspect terms (internal). The metaphysical question is already settled once we've identified true sentences which use property talk. Once we accept that properties exist, there isn't a meaningful further question about whether they *really* exist. This italicised 'really' adds no content to the question, and any sensible question in the vicinity is going to be answered by the relevant external or internal questions already identified.¹²

iv. The Easy Approach and Its Consequences

Thomasson understands questions of existence as questions about terms and their application conditions. (Thomasson 2015, 85) To understand the relationship between application conditions and existence she offers the following biconditional:

“Ks exist iff the application conditions actually associated with ‘K’ are fulfilled.” (Thomasson 2015, 86)

Thomasson claims that application conditions “are certain basic rules of use that are among those that are meaning-constitutive for that term.” (Thomasson 2015, 89) They “are among the semantic rules for use for the terms we master when we acquire language.” (Thomasson 2015, 90) She provides the following features which application conditions must have:

“1. They are semantic rules of use which speakers *master*, but these rules needn't take the form of necessary and sufficient conditions, and needn't be storable.

¹² I imagine that some may not be happy with this response. I consider it further below.

2. They are not merely conditions under which we would have *evidence* that the term applies, but rather conditions under which the term *would be correctly applied* (entitling us to truly say ‘there is a K’).
3. They need not be descriptive, and may involve deference to experts and to the world...
4. Application conditions must not take the following form: ‘K’ applies iff Ks exist. (While this will always be true, it will not count as an application condition, in our terms.)” (Thomasson 2015, 91 & 96)

Thomasson includes (4) as if we could build that biconditional into the application conditions we would need to antecedently establish that Ks exist. What the easy approach cannot require is that we have to have established existence before we can introduce a new term, as this would make the proposal circular. Application conditions form part of the rules of use for a term, and their satisfaction is sufficient to guarantee that an entity of that kind exists. (Thomasson 2015, 221)

We might worry that the entities we end up committed to from this approach may lack the features we think that they should have. (Thomasson 2015, 221) Thomasson’s solution is to state that, in addition to application conditions, concepts also require coapplication conditions. She claims:

“Application conditions are... not the only rules of use for the relevant sortal terms. It is a basic part of my view that, in addition to the application conditions, for a term to be a sortal term (or name) at all, it must also come associated with coapplication conditions: conditions determining when the term may be reapplied in a way that will entitle us to say it’s applied “to one and the same S” – thus establishing identity conditions for Ss (if any there be).” (Thomasson 2015, 222-223)

The thought is that a sortal term has two different types of rules of use. On the one hand it has application conditions which tell us the conditions under which there is something of that kind. On the other hand, it has coapplication conditions which fix the identity conditions for the things the term refers to. (Thomasson 2015, 224) Coapplication conditions “establish what we are

talking about in talking about Ss at all and establish what it means to say ‘this is the same S as that’ and what the truth conditions are for that claim.”

(Thomasson 2015, 225-226)

Thomasson considers a worry according to which there might be worlds where Ks exist but the term ‘K’ does not. (Thomasson 2015, 87) Given this concern, Thomasson opts to rigidify the application conditions for ‘K’: we stick with the application conditions for ‘K’ at the actual world. I suspect that this fix may not ameliorate all concerns in the region. Whilst rigidification will help for worlds where Ks exist, but the term ‘K’ does not, it will not help if we are inclined towards the thought that there are objects for which we have no actual associated term. Perhaps given that she is concerned with general existence questions, where we are interested in whether a certain *kind* of entity exists, e.g. properties, numbers, propositions, event, etc. she may want to hold that the only kinds of things are ones for which we already have terms. I find this option flat footed and in need of motivation. Perhaps then, more plausibly, Thomasson could follow Hale in adopting a modal account of language.¹³ Whatever form an adequate response to this concern about the expressive limitations of language takes, I proceed by assuming that Thomasson can make good on this concern somehow.

All that is required to answer ontological questions is (1) to have mastered the relevant rules for use for the term in question, and (2) to have access to the relevant conceptual and empirical information. (Thomasson 2015, 113) So: an example. We can start from an uncontroversial truth such as: (1) there are two mugs on the table. From (1) we can derive (2) analytically: (2) the number of mugs on the table is two. Thus, we are going to be committed to numbers. Thomasson therefore thinks that ontological knowledge can be gained by conceptual competence. (Thomasson 2015, 235) Sometimes this is through conceptual competence alone (in the case of an undisputed conceptual truth) and sometimes we also require empirical knowledge. The kind of conceptual competence that Thomasson wants consists in a normative claim, about what

¹³ For an indication of how this might go, see Hale (2013, 20).

speakers are *entitled* to infer, not a descriptive one about what they actually do infer. (Thomasson 2015, 239) She claims:

“Speakers who master the relevant conceptual/linguistic rules are entitled to make the relevant inference, and to accept the conceptual truth (and are open to reproach if they refuse to).” (Thomasson 2015, 240)

The relevant kind of conceptual competence is intimately tied to use.

(Thomasson 2015, 243) The thought is that a speaker must have mastery of certain core constitutive rules of use in order for those terms to be used at all. (Thomasson 2015, 250) Those may be changed, but in choosing to change the rules we do so for pragmatic reasons, not because we discover “that what seemed like conceptual truths or valid inference in the old system really weren’t.” (Thomasson 2015, 252)

Thomasson’s claim here seems to preclude the possibility that we can revise the rules for use of a term for empirical reasons. Consider: upon the discovery that space-time is curved, some things we thought to be analytic of space turned out not to be. Thus, the analytic entailments from which we derive ontological commitments would have had to change. Perhaps we should interpret such change as pragmatic, and thus we moved to a new framework where instead of using term ‘space’ we now use the term ‘space*’. This seems too radical. It seems that we can alter the meaning of terms without ‘changing the subject’. I suspect that Thomasson’s motivation in making this seemingly radical claim is that she wants to preclude the possibility that *metaphysical* considerations can force us to change the rules for use. However, I think she can and should allow that communities of language speakers can revise the rules for use for a term in light of new empirical information, without committing to the view that distinctly metaphysical considerations can force us to revise rules for use.

Given the view of ontological inquiry that is being proposed, a few comments about how Thomasson is thinking about existence is merited. Thomasson treats the quantifier as having a fixed, formal rule for use. (Thomasson 2015, 82) Her deflationary conception of ‘exists’ holds that we grasp the concept by

“grasping certain trivial platitudes governing the concept.” (Thomasson 2015, 88) Consequently, the deflationist rejects heavyweight attempts to give a substantive theory about the nature of the concept. Thomasson claims that this view is compatible with thinking that there is a relevant concept, but how we acquire this concept means that it cannot serve an *explanatory* function. (Thomasson 2015, 88-89)

In arguing against a substantive conception of existence Thomasson claims that given how it is we determine our ontological commitments, there cannot be a unified substantive or material criterion for existence. (Thomasson 2015, 117) As there are not general application conditions, so there will not be a general shared criterion of existence. (Thomasson 2015, 89) The mistake is to think that we can take the features of some kinds of existing things and generalise. This commitment in turn generates scepticism about approaches to metaphysics that propose to eliminate entities of a certain kind because they do not satisfy a certain criterion (e.g. mind-independence, etc.) (Thomasson 2015, 121) She holds that there is something wrong with metaphysical disputes which proceed by arguing that entities of a certain sort fail to meet some criterion of existence.

With a comprehensive picture of Thomasson’s proposal on the table, I now turn to consider how she thinks about the status of her proposed method for ontology. Thomasson claims that the easy approach to ontology commits one both to a first-order simple realism and to a form of metaontological deflationism. I take these commitments in turn.

Thomasson is committed to a “simple realism” about philosophically disputed entities. (Thomasson 2015, 22) If we seem to be uttering true sentences about a certain subject matter, we can then proceed as above and derive an ontological commitment to the entities in question, and we are committed to realism about those entities. Thomasson wants to resist the claim that the fact that our ontological commitments are arrived at in the way she recommends, entails that the things to which we are committed are somehow reduced in ontological standing. The things to which we are ontologically committed are not “mere shadows of language.” (Thomasson 2015, 145) We should not

understand their existence in a deflationary manner, but rather we should adopt a *simple* realism about them. (Thomasson 2015, 146) She claims that:

“The fact that we come to know of the existence of certain things by undertaking trivial inferences does not show that the entities *themselves* are in any way epistemically diminished or ontologically shallow – or that there is some crucial difference between them and regular old concreta like trees.” (Thomasson 2015, 148)

There is still a difference between trees and propositions, but Thomasson does not want such a distinction to be made along the lines of ‘ontological depth’.

(Thomasson 2015, 150) The difference is rather with respect to which truths we infer the existence of which entities (i.e. conceptual or empirical truths).

Whether we start from a conceptual or an empirical truth will likewise have some effect on how the inquiry proceeds. We may need to engage in empirical work in the case of an empirical claim, to discern what follows from what.

Whereas, “inferences made from a conceptual truth are genuinely something-from-nothing inferences, for the truth of their premises requires nothing of the empirical world.” (Thomasson 2015, 151) In the case of a conceptual truth, we should view these as articulations of the rules of use for the concept in question. (Thomasson 2015, 152-153) Thus she claims:

“Easy ontology leads to the realism about the questioned entities in affirming that there are properties, propositions, numbers, and so on, *in the only sense these terms have.*” (Thomasson 2015, 153)

However, this realism is to be distinguished from other forms of realism (e.g. Platonism) in that while all parties are signed up to the idea that the entities in question exist, the easy ontologist does not think that these entities are metaphysically explanatory. (Thomasson 2015, 155) Thomasson distances her project from what she terms ‘explanatory realism’. She claims that the easy ontologist cannot appeal to the explanatory power of the entities in question and claims that: “if an existence claim is derived by trivial inferences from an uncontroversial claim, it cannot contribute any *more* explanatory power than we got from the uncontroversial claim itself (and it can’t explain the truth of the uncontroversial claim).” (Thomasson 2015, 156) So, on the simple realist view

we can say happily that the disputed entities exist, but they should not be taken to be explanatory ‘posits’ which are part of a metaphysically explanatory theory, and our commitment to them should not be viewed as being justified by their explanatory power. (Thomasson 2015, 157)

I take it that the sense in which ontological commitments are not explanatory is subtle and open to misunderstanding, so let’s consider a toy example. Imagine two forms of ‘explanatory realism’ in mereology: compositional nihilism and compositional restrictivism.¹⁴ Let’s imagine further that all parties agree that the chair is responsible for me not falling on the floor. Now, the nihilist thinks that what *explains* why I do not fall on the floor is that there are simples arranged chair-wise. Contra this, the restrictivist thinks that what *explains* why I do not fall on the floor is that there is a composite object. These are the sort of explanatory claims that Thomasson precludes on her approach to realism. The relevant empirical explanatory claim is made clear in the statement that we start inquiry with, that the *chair* is responsible for me not falling in the floor, and there is no more metaphysically explanatory work to be done by saying that the chair is *really* simples arranged chair-wise, or *really* a composite object.

I’m sure we can imagine a hard-ontologist pushing back on this and claiming that this simple realism just isn’t what is intended by realism. Part of what it means to be a realist is to think that ontological commitments are explanatory in just the sense denied. I think that there is something of a dialectical impasse at this point, but in Thomasson’s defence it is worth considering exactly what useful sense of explanation her view fails to capture. Her approach doesn’t preclude causal explanation, scientific explanation, historical explanation, psychological explanation. What her approach rules out is that there is some distinct and weighty kind of non-empirical explanation. I suspect a simple realist would simply say that this form of explanation is not a form worth having.

¹⁴ Very minimally, I’m taking mereological nihilism to be the thesis that objects with proper parts do not exist, and mereological restrictivism to be the thesis that objects with proper parts exist.

This brings us to the sense in which Thomasson's view is deflationary. What is deflated for Thomasson is not the entities, but rather the debates about the existence of those entities. (Thomasson 2015, 154) Given that if the ontological questions are well formed, we can have answers to ontological questions, Thomasson thinks that protracted disagreement shows that something is wrong with these debates. (Thomasson 2015, 213) However, she does think that we can still make sense of some kind of disagreement between Platonists and nominalists, and this disagreement can be serious. The important thing is that we locate their disagreement as being on pragmatic rather than factual grounds. (Thomasson 2015, 175)

Thomasson considers a challenge according to which the easy approach makes ontology unacceptably language dependent. However, Thomasson claims that her deflationary view does not commit us to such language relativity. So, for example, if one part of the language (e.g. the language of fundamental physics) commits us to some kinds of entities, and another part of the language (e.g. the language of organic biology) commits us to other kinds, we don't need to say that those commitments are just relative to the relevant part of the language. Her thought is that "whether the (well-formed) noun terms appear as the basic terms of a language or as terms introduced via other parts of speech, they refer just in the case the actual application conditions for those terms are fulfilled." (Thomasson 2015, 283) So long as the application conditions are fulfilled, this is sufficient to say that the things of the relevant kind exist, and this is not dependent on the part of the language with which we are dealing.

IV. Truth for Metaontological Deflationism

v. An Important Role For Truth

The approach to ontological inquiry considered in this chapter is appealing to a philosopher who is sceptical of distinctly metaphysical methods. The

metaontological deflationist approach demystifies the methods and the epistemology of ontology. Nonetheless, there are a number of concerns that have been raised about both of the approaches examined so far. However, whilst these concerns are interesting I do not revisit them here. Rather, I will limit my discussion to the work that truth is doing in their frameworks.

Neither Hale nor Thomasson make much explicit reference to the *concept* of truth. Thomasson claims that we cannot be committed to the literal content of the utterances without also thereby being committed to the existence of the suspect entities. I take it that the reason for this is that the two statements are making the same demands on the world, and given that the statement we start with is true, the entities must exist. Whilst Thomasson's view commits her to the notion of application conditions being fulfilled (rather than satisfaction of truth conditions), the notion of analytic entailments further secures a pivotal role for truth. Hale claims that names and predicates successfully discharging their semantic functions requires that they refer. I take it that there is a broadly Fregean theory of meaning in the background of this commitment, which commits Hale to a truth conditional theory of meaning. What I take to unite their approaches is the following: truth is a feature of the sentences from which we start ontological inquiry, and by successfully semantically analysing these sentences we derive our ontological commitments.

I think that it is interesting at this point to note the similarities with Lewis's approach. Whilst Lewis offers a more sophisticated conception of semantic analysis, the idea that the semantic analysis of the true sentences of the language is what determines (up to a point) ontological commitments is a view that all parties share. The important difference, however, is that for the metaontological deflationist, semantically analysing the true sentences of the language is sufficient for fully delivering ontological commitments. There is no further metaphysical work to be done. Whereas for Lewis, there is a distinct metaphysical project that kicks in at this point, where we use metaphysical methods to discover what the nature of these commitments are.

Having identified this theoretical role for truth, I want to make explicit just how significant a role truth plays in this conception of ontological inquiry. Articulating the distinction between the different kinds of deflationism makes

this clear. Consider an interlocutor who thinks that the deflationary approach that Hale and Thomasson advocate must deflate ontology. Although all that is intended is to deflate the *project* of ontology, not the things to which we are ontologically committed, perhaps this leaves room for the question of whether metaphysics done in this way nonetheless delivers entities which are in some sense not robust. If the only way to prosecute answering ontological questions is via logical or conceptual methods (i.e. no distinct *metaphysical* method) then perhaps it follows that the deliverances of such methods are in some way language dependent or mind-dependent.

We can imagine how this charge would go. When I engage in ontological inquiry I want to know what kinds of things there are. What we can and cannot say about the world, and how our language functions, is orthogonal to this question. The intuition driving a more committal conception of ontological inquiry is that when we engage in this kind of metaphysical inquiry we want to get behind/below the level of representation, to what there is. By claiming that all there is to ontological inquiry is semantically analysing language, Hale and Thomasson fail to get behind/underneath/outside of the level of representation, and consequently cannot do justice to our metaphysical aspirations.

Hale and Thomasson both have the same answer to this worry. It is because we require the *truth* of the relevant statements from which we begin logical or conceptual inquiry that the entities to which we end up committed are not language dependent. Semantically analysing language is not alone sufficient to determine ontological commitments. Whilst our identification of certain expressions as being of a particular type, and the articulation of the application conditions for terms, is derivative on the analysis of language, the truth-values of statements containing such expressions need not be. What ensures that we only deflate the project of ontology, and not the entities, is that we require the truth of the statements from which we start semantic analysis. The motivation for assigning this robust role to truth is that it ensures some worldly connection for language. So, in response to the challenge above, the metaontological deflationist ensures that the ontological commitments

themselves are not deflated by insisting on the truth of the statements from which we begin inquiry.

vi. Three Commitments of Metaontological Deflationism

I argued that truth plays an indispensable role in the metaontological deflationist approach to ontology. It is that feature which the statements from which we start semantic analysis must have in order to guarantee that we only deflate the project of ontology and not the entities to which we are committed. Given that it plays such a robust role in the approach, it is important to articulate just what conception of truth such an approach requires. You might think that if you want to adopt the metaontological deflationist approach to ontology you can just slot in your preferred theory into the role identified. This is not so. In the course of outlining their approaches, very little is said by either Hale or Thomasson about how they are thinking about the concept. What I want to do in this section is twofold. I identify and articulate three important constraints on the concept of truth to which the metaontological deflationist must adhere, and I concurrently show that a significant number of the extant conceptions of truth cannot adhere to these constraints.

The first constraint is that whatever the concept of truth is, it must be such that it can play the robust theoretical role assigned to it. Given the role identified in the previous section, the aim of this constraint is to merely make sure that whatever we say about the concept of truth, it is capable of playing this role. I take it that this constraint is sufficient to rule out a redundancy theory of truth. There is a much subtler matter about the degree to which this constraint is compatible with deflationism about truth. Hale and Thomasson have both expressed their preference for a deflationary conception of truth.¹⁵

Whether this is a tenable position will depend in large part on how the relevant kind of deflationism is understood. It is worth noting that there is very little philosophical consensus about just what deflationism about truth commits you

¹⁵ For example, see Thomasson (2014).

to. You might think that a ‘deflationary conception of truth’ can be explicated along the following lines: the cognitive equivalence of “P” and “‘P’ is true”, minimal introduction and elimination rules, perhaps some other more minimal Wrightian platitudes.¹⁶ Further, it seems plausible that this way of thinking about deflationism is compatible with thinking that truth is explanatorily potent (at least inasmuch as the role identified here demands). I will adopt something like this conception of deflationism in the following section. However, it is worth noting that there are other conceptions of deflationism which argue for deflationism by drawing attention to the putative explanatory impotence of the concept. I take it that if this is the type of deflationism that you favour, then this type will be ruled out by the first constraint.

The second constraint is that whatever the concept of truth is, it must not yield an anti-realist metaphysics. You might think that this is too strong. Take a standard characterisation of realism to be a commitment to two theses: a claim about existence (i.e. objects of the relevant kind exist) and a claim about independence (i.e. the existence of those objects is not dependent on human minds, languages, theories, etc.). On the metaontological deflationist approach, existence is secured easily. What of independence? Both Hale and Thomasson explicitly commit themselves to views which appear to require independence: part of the Neo-Fregean project is to show how we can embrace Mathematical Platonism, and Thomasson aims at a ‘Simple Realism’. I suspect that the drive behind their commitment is that in the face of challenges to their approach they want to reject certain anti-realist platitudes. Hale and Thomasson reject the view that ontological commitments derived in the way they suggest: (1) depend for their existence on human language or minds, (2) only exist because we have a language to describe them, (3) would cease to exist if we did, etc.

Whether this is a robust enough conception of ‘independence’ to count as a *really* realist position, I leave to the reader to decide. What is important for my present purposes is that given this constraint, some epistemically constrained conceptions of truth will not do. Again, explicating what the commitments of

¹⁶ See Wright (1992).

epistemically constrained conceptions of truth are requires subtly. There certainly are, however, conceptions of epistemically constrained truth which will be ruled out. For example, this constraint will rule out some conceptions of truth where, for example, being true is analysed as holding true. I remain silent on whether there are more sophisticated forms of alethic anti-realism which are compatible with metaontological deflationism, but radical constructivist positions are ruled out.

The third constraint is that whatever the concept of truth is, it must not yield an explanatory metaphysics. The ontological commitments generated by the approach should not be taken to be explanatory ‘posits’ which are part of a metaphysically explanatory theory. Thomasson says: we do not get to say that our ontological commitments “*explain* what it is that makes our sentences true.” (Thomasson 2015, 157) We need to be very careful with this constraint. I think that the commitment motivating this constraint is the following: it must be the case that we are permitted to *start* from the truth sentences of the language in order to determine the ontological commitments. Think about the fourth feature that Thomasson claims application conditions must have:

“Application conditions must not take the following form: ‘K’ applies iff Ks exist. (While this will always be true, it will not count as an application condition, in our terms.)” (Thomasson 2015, 96)

What we need to be able to do if we are metaontological deflationists is take legitimate uses of the truth predicate to not require antecedent metaphysical assurance. A consequence of this way of thinking about truth requires that we don’t want a conception of truth where we will be able to go back and consequently say that the statements we thought were true are in fact false. Imagine the following. We start with some class of statements that we take to true, say moral discourse. We semantically analyse the statements and determine that the truth of these statements requires commitment to moral properties. Now, what cannot happen for the metaontological deflationist is that we discover that actually the world contains no such properties and so the class of statements we started with is really false. There is a general rejection of error theories for the metaontological deflationist.

I can imagine some resistance to this claim. Consider phlogiston. It seems like we have a case of a meaningful term, which at one time appeared in statements which were thought to be true. Then, scientists discovered that there was actually no such thing as phlogiston, and so the statements that were previously believed to be true, were in fact false. Isn't this an acceptable case of error theory? I suspect that the metaontological deflationist needs to say no. What happened in the case of 'error theory' about phlogiston is that science progressed and this altered the class of statements from which we derived ontological commitments. It was not by metaphysical reflection on the nature of the world that we arrived at the conclusion that phlogiston does not exist, and thus the phlogiston involving statements are false.

I think that Reichenbach's distinction between the direction of discovery and the direction of justification can help to illuminate this more. (Reichenbach 1938) The direction of epistemological explanation and the direction of metaphysical explanation do not have to run in the same direction. For the metaontological deflationist, given that they are deflating the project of ontology, the only direction of explanation with which they are concerned (inasmuch as metaphysics is concerned) is the epistemological one. We have no access to facts about the direction of metaphysical explanation, and so talk of ontological commitments explaining the truth of statements with which we start cannot do justice to the desired deflationism. Given this constraint, conceptions of truth in terms of truthmaking are ruled out cannot do justice to this constraint.

vii. Truth for Metaontological Deflationists

I have argued that truth plays an indispensable role in the metaontological deflationist framework, and that there are three constraints on how the operant notion of truth can be understood. I argued that this rules out a number of approaches to truth. In this section I offer a constructive proposal for how the metaontological deflationist could articulate the conception of truth that they need.

Before I start it is important to acknowledge a dialectical tension. I contend that we aren't going to be able to get a handle on the right way to think about truth in the metaontological deflationist framework by engaging with questions about *what truth is*. If the metaontological deflationist aims to deflate the project of ontology by showing how we needn't use metaphysical methods to answer ontological questions, then approaching the question of what the right conception of truth is for this project by asking questions about the nature of truth which cannot be answered *easily* would undermine the very aspirations of the metaontological deflationist project. It is therefore my contention that the metaontological deflationist should rather look to the theoretical role assigned to truth in the theory of meaning and interpretation to articulate the notion that they require. In what follows I make a suggestion for an approach to thinking about truth that I take to be fruitful for the metaontological deflationist.

Consider Davidson's metasemantic project of radical interpretation that I outlined in Chapter 4. Recall that a methodological assumption on the behalf of the interpreter is the principle of charity. Davidson claims that speakers hold "a sentence to be true because of what the sentence (in his language) means, and because of what he believes." (Davidson 1973b, 134) We saw that charity can be thought of a commitment to two further principles: coherence and correspondence. Importantly for my present purposes, correspondence says that the interpreter should take there to be a relationship between the beliefs and utterances of speakers of a language and their environment. Thus, attributions of belief and assignments of meaning must be consistent with the evidence afforded by our knowledge of the speaker's environment.

The goal of the interpreter is to generate a T-sentence for every sentence of the object language, whilst respecting the principle of charity, and radical interpretation proceeds by holding belief fixed to work out meaning.

(Davidson 1973b, 137) On such an approach, truth then gets into the picture by playing an explanatory role in the theory of interpretation: it is that feature which we take sentences to have when we assert or believe them. Further, because of the methodological presuppositions for radical interpretation, we can, if we like, talk about the content of belief or sentences as corresponding to

the way the world is. Built into the principle of charity was the claim that there is a relationship between the content of beliefs and utterances and speakers' environment, and, in negotiating an interpretation to make allowances for explicable error, the interpreter will bring in environmental factors as explanans. Thus, when we use the theory of truth to give the theory of meaning in Davidsonian style, and we engage in a radical interpretation project, this commitment results in the empirical claim that there is a correspondence between the content of our utterances and the worldly things they are about. Importantly, given my purposes in this chapter, the motivation is not hefty metaphysical considerations to do with the nature of truth, but rather naturalistically respectable constraints on the theory of interpretation.

Returning to the constraints that I articulated in the last section, we can see how the way of thinking about truth that I advocate here conforms to these constraints. First, this 'interpretationist' conception is capable of playing a robust theoretical role. In addition to this, because of the principle of correspondence it gives us something principled to say about why and how it can play this role. Second, it doesn't furnish us with an anti-realist view, at least in principle. Whether this way of thinking about truth requires alethic anti-realism will in part depend on what you think the right notion of truth is for the theory of interpretation. So, for example, if you are convinced by Dummettian arguments that the right way to understand the conception of truth at work in the theory of interpretation is such that it is not bivalent, then perhaps the way of thinking about truth that I advance here will not deliver alethic realism. However, given that the status of the metaontological deflationist independence condition is in need of further articulation, I am confident that a consistent picture can be constructed. Third, and finally, my recommendation articulates a way of thinking about truth which does justice to the sense in which metaontological deflationism requires that ontological commitments are not metaphysically explanatory. The interpretationist strategy pursued recovers some notion of correspondence, but without committing us to the short of heavy duty truthmaking that the metaontological deflationist wishes to avoid.

V. Conclusion

I have argued that if the metaontological deflationist approach to ontological inquiry is to be workable, it needs to be furnished with an adequate account of truth. Further, I claimed that as it stands, an adequate account of truth has not been sufficiently articulated. Given the significant theoretical role that truth plays, this should be worrying. On the one hand, making good on the epistemological and methodological virtues of the approach is held hostage to the development of such an account. On the other hand, and more importantly for my current purposes, for the metaontological deflationist's conception of the role of truth in metaphysics to be compelling, they need to be able to provide some kind of story about how it is that truth can play the role assigned to it. Nonetheless, I suggested that an adequate notion could be articulated, and I have gestured towards an approach I believe to be a fruitful in attempting to do just that. Of course, the acceptability of this way of thinking about truth is going to depend on whether it is an acceptable way of thinking about the role of truth in the theory of meaning.

Chapter Six: Michael Dummett, Hilary Putnam, and The Metaphysical Consequences of Epistemically Constrained Truth

I. Introduction

This chapter investigates how commitment to an epistemically constrained conception of truth can affect the possibility and nature of metaphysical inquiry. I consider two historically influential arguments for the claim that truth is epistemically constrained: one from Dummett, according to which we must hold that truth is epistemically constrained once we do justice to the role that truth plays in the theory of meaning; and one from Putnam, according to which considerations from model theory show us that truth cannot be such that it is epistemically unconstrained. Like Lewis, Dummett and Putnam each aspire to a univocal conception of truth: whatever truth concept plays a role in metaphysical inquiry, it is identical to the one that plays a role in the theory of meaning. Contra Lewis, however, both Dummett and Putnam take it to be the case that once we articulate the demands placed on the concept by the theory of meaning, this requires that we reject a certain conception of the metaphysical enterprise. Whilst different sets of considerations motivate Dummett's and Putnam's respective arguments for epistemically constrained truth, they both take this alethic commitment to preclude the possibility of realist metaphysics.

Before commencing the discussion, it is useful to provide some provisional clarifications about two salient notions in this chapter. First: epistemically constrained truth. Neither Dummett nor Putnam *define* truth in terms of some epistemic notion or capacity. The predicate 'is true' is not defined as the predicate (e.g.) 'is warrantably assertible'. I dedicate some time to elucidating this distance between Dummett's project and definability in the next section.

In addition to this point about definability, it is important to register that neither philosopher holds the obviously false belief that all truths are known. Rather they reject what I'll term alethic realism (following Künne 2003, 20): the view that some truths are evidence transcendent. The alethic realist thinks that for domains which aren't effectively decidable there could be a kind of inevitable ignorance; that there are some sentences which human beings are capable of understanding, but which could never be the content of a justified belief. Those who hold that truth is epistemically constrained deny this and claim that truth cannot outrun what it is in principle possible for some human (i.e. cognitively finite) being to be justified in believing. Dummett and Putnam argue against slightly different manifestations of alethic realism: Dummett argues that truth is not bivalent, and Putnam argues against the view that there could be undetectable error.

By way of a second clarification, it is useful to make precise the sense in which Dummett's and Putnam's arguments impact upon the status and nature of metaphysical inquiry. In both instances, their discussion of metaphysics does not concern first-order metaphysical theses, but rather pertains to what might be called second-order metaphysics: the disagreement between realism and anti-realism. This is a divergence from the sort of issues which have preoccupied previous chapters. Whilst all the other philosophers and approaches I have considered offer a particular method which makes use of their particular conception of truth, Dummett's and Putnam's programs aim to rule out the possibility of a certain kind of metaphysics directly. So, for example, the metaphysical method advocated by Thomasson further commits her to a type of metaontological deflationism. Likewise, the method offered by Davidson further commits him to a type of structuralism. They offer a particular method for answering metaphysical questions and this method has implications for the subsequent status of metaphysics. Dummett and Putnam, however, each present specific arguments against metaphysics having a particular status.

The chapter proceeds as follows. First, I consider Dummett. (§II) I outline his theory of meaning, the central role assigned to the concept of truth in this project, and how this role constrains the nature of the property which can play

this role. In turn, I explain how this commitment precludes evidence transcendent truths, and requires the rejection of realism. Second, I consider Putnam. (§III) I briefly trace the development of his thought about truth from the Putnam we encountered in Chapters Two and Three to the Putnam of the model theoretic argument. I outline the model theoretic argument, and how this argument aims to establish both the impossibility of realist metaphysics and the epistemically constrained nature of truth. Finally, I close this chapter by making some general remarks on the relationship between metaphysics and epistemically constrained truth. (§IV)

The aim of this chapter is not to establish that truth is in fact epistemically constrained. The aim of this chapter is to evaluate the supposed connection between truth and realism. Dummett and Putnam each think that it is a constitutive thesis of realism that truth is not epistemically constrained. So, when they show that truth is epistemically constrained, we have to reject metaphysical realism. One way of resisting this argument is to show that truth is epistemically unconstrained. Another route is to undermine the connection between truth and realism. This is the approach that I investigate here. In particular, I want to defend the following conditional: if truth is epistemically constrained, then this has unavoidable metametaphysical consequences. Defending this conditional requires some significant revision of what Dummett and Putnam take to be the realist/anti-realist consequences of their respective arguments.

II. Dummett

The purpose of this section is to outline Dummett's case for epistemically constrained truth and explain how this commitment rules out a certain kind of metaphysics. Dummett provides a clear statement of what he takes to be the insight of his position:

“The whole point of my approach to these problems has been to show that the theory of meaning underlies metaphysics.” (Dummett 1978, xl)

This section attempts to unpack this claim. There are two interconnected components: a conception of the theory of meaning, and a view about the nature of metaphysics. I claim that truth, for Dummett, is operationally constrained by the role that it must play in the theory of meaning, which is in turn constrained by the nature of linguistic understanding. That truth is thus constrained precludes a realist metaphysical enterprise.

i. Truth and Meaning

There are two clear influences on Dummett's thinking about truth and meaning, which he explicitly acknowledges: Frege and Wittgenstein¹⁷.

Dummett thinks that the central notions for a theory of meaning are going to be notions of truth, truth-conditions, and compositionality.¹⁸ How these notions should be foundationally understood is going to turn on facts about language acquisition and manifestation. So we must offer a theory of meaning which has a formal structure such that we can explain how the truth values of syntactically complex sentences depends on the truth values of the composite parts, but the relevant notion of truth must be constrained by facts about how we use language, how we make and justify assertion, etc. Dummett therefore proposes a theory of meaning which consists in giving epistemically constrained truth conditions. In this subsection I explain what this meaning theory is, and what motivates it.

It is worth noting, in the interests of clarity, that Dummett's thinking about the role that truth should play in the theory of meaning changed quite substantially. There is a popular reading of Dummett according to which he eschews the notion of truth in the theory of meaning altogether. In its place, he rather favours a theory of meaning in terms of assertibility conditions.

(Williamson 2014, 22) Whilst it is correct that at one point Dummett did

¹⁷ Although he emphasizes that post-1960, he wouldn't have identified as a 'Wittgensteinian'. (Dummett 1978, xii)

¹⁸ Throughout this section where I talk about Dummett's theory of meaning I am focussing on his theory of *sense*.

favour an assertibility conditional theory of meaning, by *Truth and Other Enigmas* the mature Dummettian position is that the meaning of a sentence should be given in terms of its truth conditions. Thus, what is revisionary about this mature Dummettian position is not the rejection of a truth conditional theory of meaning, but rather a substantial revision of what the permissible conception of truth must be.

Dummett is signed up to the idea that linguistic understanding consists in knowledge of the conditions under which a statement would be true. Whilst he takes it as philosophical orthodoxy that this relationship between meaning and truth conditions holds, he claims that the nature of this relationship is often left insufficiently elucidated. Specifically, he thinks that those philosophers who think that merely offering something like a Tarskian truth definition¹⁹ is sufficient for capturing the conceptual and theoretical role of truth are mistaken. (Dummett 1978, xxi) As he says: “the problem is not whether meaning is to be explained in terms of truth-conditions, but of what notion of truth is admissible.” (Dummett 1978, xxii) What’s wrong with the Tarski definition according to Dummett is that it doesn’t do justice to the theoretical role of truth. In fairness to Tarski, the explicit definition he gives of truth in terms of satisfaction by all sequences isn’t intended to have the sort of theoretical significance which Dummett takes an adequate elucidation of the concept to require. Dummett thus accepts the T-schema, but thinks that further elucidation of the concepts of truth is required.

Whilst truth plays a central and indispensable role in his semantic theorising, Dummett does not think that it is a fundamental notion. (Dummett 1978, xxiii) Satisfactorily explicating Dummett’s claim that truth is not a fundamental notion requires subtlety. Dummett is committed to the Fregean thesis that truth is somehow conceptually basic. In different places Frege claims that truth is irreducible, indefinable, and not explainable. (Künne 2003, 16) In trying to articulate Dummett’s precise view about the relative fundamentality of the concept of truth by pointing to his Fregean commitment, it would be messy of me to seamlessly elide reduction, definition, and explanation. All of these

¹⁹ See Tarski (1936).

claims say slightly different things about truth, and it is important to distinguish them. I suggest therefore that we should interpret the Fregean thesis as a claim about definition: truth cannot be conceptually analysed into more basic concepts. However, in addition to this claim that truth cannot be defined, Frege also holds that “in the laws of truth the meaning of the word “true” is spelled out.” (Frege, quoted in Künne 2003, 17-18) So, whilst the concept of truth cannot be analysed, we can say informative things about the concept by looking to the role it plays in logic: truth is that which is preserved by valid deductions.

There is of course much more to say with regards to these exegetical subtleties in Frege. However, this is enough to get a handle on Dummett’s view. When Dummett claims that truth is not fundamental we should take this claim to pertain to explanatory rather than conceptual priority. That is: to get a handle on the kind of concept truth is, we do not offer an account of it in more basic terms, but rather we must look at the theoretical role that it plays. Dummett thinks that the primary role for the concept of truth is in the meaning theory.

So, what role does truth play in the meaning theory? Dummett holds that the meaning of sentences of our language can only be given when there is a method for determining their truth values that we are in principle able to carry out. The relevant notion of truth for Dummett is one which is to be explained in terms of the success of assertion. (Dummett 1978, xvii) Truth and falsity are therefore inextricably tied to our linguistic and epistemic practices. Dummett urges that to explain why this is so we need to start from a question about what the meaning theory is supposed to do. (Dummett 1993, 16) He thinks that an answer to this question is given by reflecting on the nature of linguistic understanding. He claims:

“A semantic theory is not a complete meaning-theory but only a preliminary outline sketch of one; and it cannot be judged correct or incorrect until it has been expanded into a meaning-theory which displays the connection between the meanings of the sentences, as represented by the theory, and the practice of using the language.”
(Dummett 1993, 18)

Dummett offers two arguments to the effect that truth must be epistemically constrained, both of which trade on claims about what is required for linguistic understanding. First: the acquisition argument. As I have said, Dummett is committed to the Fregean thesis that to understand a statement, we must know its truth conditions. Dummett holds the *prima facie* plausible constraint on knowledge ascriptions that in order for knowledge ascriptions to be legitimate it must at least in principle be possible for that speaker to have *acquired* that knowledge. Since we could not acquire knowledge of evidence transcendent truth conditions, truth conditions must in principle be knowable. Second: the manifestation argument. Starting from the same Fregean assumption, Dummett adds to this the Wittgensteinian thought that understanding does not consist some kind of inner state, but rather consists in a kind of practical ability. Thus, if we ascribe knowledge of truth conditions to a speaker, it must in principle be possible for that speaker to *manifest* that knowledge in their use of the language. Knowledge of evidence transcendent truth conditions cannot be manifested in our use of a language. So, again, given these considerations to do with linguistic understanding, truth conditions must in principle be knowable.

The claim that truth conditions must be in principle knowable requires exegetical subtlety. It is clear there is some kind of verificationist commitment here and so the important question is how much verificationism. Following Putnam, Dummett concedes that:

“It is misleading to concentrate too heavily, as I have usually done, on a form of anti-realist theory of meaning in which the meaning of a statement is given in terms of what conclusively verifies it; often such verification is not to be had”. (Dummett 1978, xxxviii)

In this spirit, then, a mature Dummett rejects what he describes as extreme constructivism:

“[to deny] that there are true statements whose truth we do not at present recognise and shall not in fact ever recognise... would appear to espouse a constructivism altogether too extreme. One surely cannot equate truth with being recognised... as true”. (Dummett 1993, 446)

In these passages Dummett wants to resist the thought that epistemically constrained truth requires that all truths are in fact known.

Let's consider an example. Take two arbitrary numbers, say 3573624 and 3695827. These two numbers have a product, but let's imagine that no one has ever actually multiplied these numbers together before. Thus, there is a sentence, 'the product of 3573624 and 3695827 is x' of which no-one knows the truth value. Further, is it possible that no-one may ever know the truth value of this sentence. However, the reason for this is merely that no one bothered to find out. If someone wanted to find out whether this sentence is true, there exists a method for determining its truth. Contrast this with the sentence 'every even integer greater than 2 can be expressed as the sum of two primes'. The explanation of our ignorance about the truth of this claim seems significantly different to the previous claim. There isn't a procedure for ascertaining the truth of this sentence. All of this is to say that there is a significant distinction between the truth of a sentence being known and it being knowable.

So, understanding a particular statement involves knowing its truth condition. We require knowledge of truth conditions to do compositional semantics. But what it is to know the truth condition for that statement is to know whether it could be correctly asserted or correctly denied. The modal distinction (i.e. between the truth of a sentence being known or knowable) needs to be spelled out in terms of verification procedures. To do this, without falling into the sort of extreme constructivism that Dummett rejects, we need to use counterfactuals.²⁰ The thought is that if we have actual or counterfactual procedures for ascertaining the truth of the sentence under consideration then we can say that the statement is either correctly asserted or correctly denied. Where there exists no procedure for determining the truth value for a statement, even counterfactually, we must refuse to assert that the statement has a truth condition. Consequently, the statement in question cannot be said to be determinately true or determinately false. The requirement that all truth

²⁰ This isn't to give an *analysis* of verification procedures in terms of counterfactuals, but rather to just help elucidate the notion. See Dummett (1993, 215-247) for details.

must be knowable proves to have consequences for Dummett's view of which logical laws to accept, which, in turn, has consequences for the metaphysical question of realism. The next section is devoted to elucidating these commitments.

ii. Truth and Realism

Dummett claims that the debate between realism and anti-realism as it has been traditionally conceived is misguided. In the Preface to *Truth and Other Enigmas* he claims that when considering a certain domain of inquiry we have two questions to ponder. First, we might wonder how we settle a dispute over the ontological status of the objects in question. Second, we might wonder what bearing either metaphysical view (realist or anti-realist) has on the forms of reasoning employed within the debate. (Dummett 1978, xxv-xxvi) In posing these two questions, we start to get an idea of the fact that Dummett thinks the questions of metaphysics and the questions of what logical laws to accept are importantly connected. Indeed, he claims that as they normally proceed, debates about realism obscure that there is a single issue at stake:

“There could, indeed, be said to be a single higher-order issue: under what circumstances are we entitled to assume the principle of bivalence for some class of statements?” (Dummett 1978, xxxi)

It is important to distinguish the two distinct components of Dummett's views about metaphysics. There are ontological questions: questions about what kinds of things exist. The Fregeanism in Dummett allies him with the neo-Fregean view explored in the previous chapter, according to which ontological questions are in some sense easy to answer. We are committed to a candidate object existing when the corresponding singular term occupies a reference demanding position in a true sentence (and the same *mutatis mutandis* for predicates and properties). There is thus no further ontological question to be answered once we know which statements are true. The other dimension of metaphysics is concerned with objectivity of the subject matter. Whether the objects to which we are committed are 'independent'. Dummett holds that

independence is not “a possible object of discovery alongside the propositions it proposes to interpret: it is a doctrine concerning the status of those propositions.” (Dummett 1993, 8) So how should this independence thesis be understood? Dummett thinks that the non-metaphorical understanding of independence is bivalence. Realism is thus concerned with whether the statements pertaining to a particular ontological debate are determinately true or determinately false. (Dummett 1978, xxxii)

We might wonder why we should hold that the fight between realism and anti-realism should be over whether truth is bivalent. Dummett conceives of two ways in which we might try to tackle the question of realism. We might approach it from the top down; that is, we might proceed by attempting to deliver an answer to the metaphysical question first, and from there discern the correct conception of truth and meaning, and “hence to deduce the logic we ought to accept as governing them”. (Dummett 1993, 12) Dummett thinks, however, that such an approach cannot be viable. He thinks it is methodologically opaque how we would even start trying to answer the metaphysical questions. In particular, that evaluating the metaphysical arguments independently of considerations to do with the appropriate notions of meaning and truth, cannot help but be metaphorical. Dummett claims that conceptions of realism which hold that what is at issue is what is ‘basic’, ‘grounded’, ‘fundamental’, etc. makes the debate hopelessly obscure. Instead of this top down method, Dummett prefers a method that starts with the question of the correct model of meaning for the disputed class. Once we have ascertained the correct model of meaning for the language (or for a part of the language), we are in a position to elucidate the concept of truth in play for that language (or part of it). (Dummett 1993, 14; 1978, 146) Thus, for Dummett if we want the debate between the realist and the anti-realist to be tractable, we need to make sense of what is at issue by focussing on the appropriate notion of truth for the language. (Dummett 1978, xix)

The issue of realism for Dummett, therefore, is an issue of whether the truth values of the statements are fully determinate. (Dummett 1978, xxix) That is: does bivalence hold? Importantly, the method we use to answer this question must take as its starting point the theory of meaning:

“Thus the answer to the question how is it possible to call a basic logical law in doubt is that, underlying the disagreement about logic, there is a yet more fundamental disagreement about the correct model of meaning, that is, about what we should regard as constituting an understanding of a statement.” (Dummett 1993, 17)

Why might you think bivalence holds for a particular domain of discourse? You might think it holds because there is an effective decision procedure or an effective method of verifiability for determining whether a statement is true or false. Statements have a determinate truth value in virtue of their having some effective method for determining their truth value. The conception of truth relevant is thoroughly epistemic constrained. By contrast, you might think that truth transcends available evidence; that statements are either true or false independently of our ability to recognise them as such. In sum, we get bivalence one of two ways: (1) because there is evidence transcendent truth, or (2) because we have an effective method for ascertaining truth. Accordingly, the argument between the realist and the anti-realist is a dispute about whether all truths are knowable.

To settle whether truth is bivalent, therefore, we have to determine whether truth can be evidence transcendent. If you were of the persuasion where you thought that truth is epistemically unconstrained, it would be a consequence of this concept that truth can outstrip what is knowable. However, Dummett of course thinks that we should reject the view that there can be evidence transcendent truth, as this cannot do the theoretical work we require of the concept of truth in an empirical theory of understanding. What we learn from the manifestation and acquisition arguments is that the proper analysis of the concept of truth delivers an implicit definition where it is an analytic truth that what is (actually or counterfactually) warrantably assertible is true. Where we have no idea of what would count as warrant, talk of determinate truth or falsity does not make sense. The anti-realist does not have to come up with examples where bivalence fails – if there is no gap between warranted assertibility and truth then it is possible that it does and therefore we are entitled to reject bivalence. This way of understanding the dispute makes it a modal dispute about the consequences of the concept of truth. The anti-realist

maintains that there shouldn't be a gap between the concept of truth and the laws of truth – hence bivalence does not follow.

The law of bivalence says that for any sentence of the language, that sentence is either true or it is false. In light of the manifestation and acquisition arguments, Dummett thinks we must reject the law of bivalence, and that this will settle the metaphysical controversies over realism “without residue”. (Dummett 1993, 14) Thus, Dummett is committed to a global form of anti-realism once we realise that bivalence must go. This global anti-realism claim requires some care. In some ways, Dummett represents the methodology of choosing a side in the realist / anti-realist debate as a piecemeal approach. We must work through the individual domains of inquiry and determine whether, for this specific field of inquiry, it is such that the statements are determinately true or determinately false. (Dummett 1978, xxxix) Now, it may be the case that as a matter of fact there is some restricted domain of discourse for which every statement is determinately true or determinately false. However, the reason why this is so is because all the sentences in this restricted domain have some effective method for determining their truth or falsity, not because their truth is evidence transcendent. So, Dummett is a global anti-realist in the following sense: (1) he denies the law of bivalence and (2) if for some restricted class of sentences all those sentences are determinately true or determinately false, it is only so because all the sentences have an effective method for determining their truth or falsity.

In sum, given that Dummett is in the business of compositional truth-theoretic semantics, we get to metaphysics in some sense merely by the disquotational property of truth. However, Dummett forwards a sceptical position about classical semantics. Once we take the manifestation and acquisition arguments seriously, we see that there could not be evidence transcendent truth, which means we need to reject bivalence, and this means we need to reject realism.

III. Putnam

In Chapters Two and Three when considering what motivates the talk of truth in Dyke and Sider's respective projects, I suggested a Putnamian conception of truth. I claimed that the Putnam of *Meaning and the Moral Sciences* forwards a kind of empirical realism about truth. According to this conception, we essentially have a correspondence theory, but what Putnam deems an empirically satisfactory one. He claims that the primary constraint on what counts as a good truth-definition for a language is that it should be the one which enters into the best explanation of the behaviour of a speaker. (Putnam 1978, 41) Rationalising the behaviour of a speaker is to offer an explanation, and in order to offer this explanation we need to have an adequate handle on the connections between understanding, truth conditions, belief, action, etc. (Putnam 1978, 41) This is how correspondence gets into the picture: "Correspondence' between words and sets of things... can be viewed as part of an *explanatory model* of the speakers' collective behaviour." (Putnam 1978, 123) Putnam thinks that the fact that there is such correspondence is what explains the success of what we can do with language, and the success of natural science. (Putnam 1978, 111)

The Putnam that I discuss in this section offers a different view of the role of truth in meaning and metaphysics. Putnam's views about truth and metaphysics changed (more than once, but we'll only be concerned with the first of his transformations), and they did so in part in response to Dummett's insights. (Putnam 1978, viii) Putnam thinks that we must take seriously the thought that truth is epistemically constrained. The biggest difference this change makes is in how Putnam understands the gap between warranted assertability and truth. In the earlier work Putnam holds that there must be a difference as something can be warrantably assertible but not true. (Putnam 1978, 108) However, he later argues that we must reject this commitment and hence the traditional realist distinction between true and ideal theory.

The rest of this section is dedicated to properly elucidating Putnam's view of truth and realism. Before I do this, however, I think it is interesting to draw out a similarity between earlier period Putnam and the version of his views about truth I consider here. In *Meaning and the Moral Sciences* Putnam makes a general and constraining claim about truth. Putnam holds that how we understand the

concept of truth should be determined not merely by its formal properties, but by the role that it plays within total theory. (Putnam 1978, 37) The significant divergence between these different time slices of Putnam, then, isn't a methodological one. To give an adequate account of the concept of truth we should look to the role that it plays. The divergence comes as Putnam forwards an argument which shows that truth cannot play a role that he previously assigned to it. It is to this that I now turn.

iii. Putnam's Attack on Realist Truth

Putnam claims that his model theoretic argument "has profound implications for the great metaphysical dispute about realism which has always been the central dispute in the philosophy of language." (Putnam 1980, 464) Putnam's model theoretic argument aims to show that given plausible considerations from model theory, we are misguided in believing that our language can be hooked up with a mind-independent world in the right way; where 'right way' is to be understood as 'the way the metaphysical realist purports they are hooked up'. Putnam thinks this motivates a move away from metaphysical realism towards his preferred 'internal realism'.

Before turning to the argument itself, I need to outline how Putnam characterises the commitments of metaphysical realism. On his view, the realist is not committed to one thesis but three interrelated theses. First, a claim about the world: "THE WORLD" is to be characterised as independent, in the sense of independent of any representation of it. Second, a claim about language: that the reference relation for the realist is one of correspondence. That is: that there is a determinate reference relation which holds between expressions of our language and the parts of the world to which they refer. Third, a uniqueness thesis: that there is a *singular* correspondence between the terms of a theory and the objects and the properties to which they refer. (Putnam 1981, 49)

Expressing the uniqueness thesis requires some subtlety. Putnam claims that for the realist there is a singular correspondence between theory and the world.

We might think this is too strong a requirement and as such Putnam is overstating the realist commitments. There are some forms of semantic indeterminacy that do not seem, *prima facie*, to be a problem for metaphysical realism. There is a prevalent phenomenon of indeterminacy in natural language: vagueness. It is not particularly controversial to claim that natural language is vague: reference is not always determinate. Paradigm terms that exhibit this indeterminacy would be, 'tall', 'heap', 'bald', etc. Such terms admit borderline cases; cases where it is not clear whether the predicate applies. (Keefe 2000, 6) Such terms apparently lack clear extensions²¹. Given that there are vague expressions in a language, it seems that we need to make sense of the reference of a term not being wholly determinate. In offering an explanation of this phenomenon, a semantic approach has been historically most popular: theorists have tried to explain this vagueness in terms of features of the language, and have offered a semantics and a logic for the suspect expressions. (Merricks 2001, 146) We might think, therefore, that the existence of vague predicates suggests that a one-to-one correspondence is too committal. Consider your total theory of reality. Perhaps you quantify over colours. There could be two properties, red and red*, which the word "red" in the theory could pick out. If our best total theory contains the predicate "is red", and yet there are two equally good candidate referents for "red", this might seem like enough to suggest that there is not a singular correspondence, but rather that there may be multiple equally good ones.

I think that the realist could make three moves. First, perhaps she could retreat. A one-to-one correspondence is too committal, and as such the realist accepts that in some instances there is not a singular correspondence. The theory can correspond to greater and lesser degrees, and the realist prefers those theories which correspond more. However, there can still be a degree of flexibility, especially in those cases of mundane indeterminacy considered.²²

²¹ It is worth noting that the view that extensions are not determinate only fits with some explanations of vagueness. For example, according to Williamson's epistemicism about vagueness (1994), the referents of vague terms are determinate; the source of the vagueness is in our own ignorance.

²² Perhaps this could be metaphysically spelled out in terms of a commitment to ontic vagueness. This would not be to move away from realism as we can be realist about what is

The second option is stricter. This type of realist may object to the putative problem, and they will say that no metaphysical theory would quantify over colours simpliciter. Perhaps this is because they think that colours are not the types of things that belong in the fundamental description of reality, or perhaps it is because the language we use to talk about colours in metaphysics do not contain the predicates of everyday English but rather precisified technical predicates.²³ Thus, this realist maintains the one-to-one correspondence, by eradicating instances of mundane indeterminacy by some means. Third, perhaps the realist could advance the view of vagueness advocated by Williamson, according to which there is uniqueness, the terms in question do determinately correspond, but rather locate the source of vagueness in our own ignorance. (Williamson 1994) Putnam seems to be levelling his argument against realists of the second and third variety. For our purposes, what is important is the thought that there should be an intended interpretation for our theory, even if in some cases we cannot know what it is. The realist contention is that the intended interpretation is the one that matches the terms of the theory to the way the world is.

And so, we return to Putnam's characterisation of realism. Once we have these three theses, it should be clear that it is possible that we may be unable to represent THE WORLD at all. To put it in Putnam's terms: truth is, for the metaphysical realist, radically non-epistemic. (Putnam 1977, 485) The theory which is epistemically ideal, in the sense of meeting all our theoretical virtues, might still be false. Of course, it might be the case that ideal theory *is* true; the claim is not that ideal theory *must* fall short of truth. Rather the claim is that, for the realist, idealness does not constitute or guarantee truth. The realist must hold that there is a gap between ideal theory and true theory. At the heart of Putnam's attack, therefore, is a distinction that the realist must be committed to, but, according to Putnam, cannot maintain:

metaphysically indeterminate. For some illuminating discussion please see Barnes and Williams (2011).

²³ This approach obviously resembles Sider's move to Ontologese that I considered in Chapter Three.

“Here again, the realist – or, at least, the hard-core metaphysical realist – wishes it to be the case that *truth* and *rational acceptability* should be *independent* notions. He wishes it to be the case that what, e.g., electrons are should be distinct (and possibly different from) from what we believe them to be or even what we would believe them to be given the best experiments and the epistemically best theory. Once again, the realist – the hard-core metaphysical realist – holds that our intentions single out “the” model, and that our beliefs are then either true or false in “the” model *whether we can find out their truth values or not.*” (Putnam 1980, 472)

Now, let’s turn to the argument. I should note that “*the* model theoretic argument” doesn’t strictly speaking pick any one thing out. Putnam forwards several arguments which are similar in spirit, and, whilst they differ in precise target and technical machinery, they all aim to forward the same conclusion²⁴. Some of his discussion focuses on mathematical language and employs the Löwenheim-Skolem theorem²⁵, whereas other discussion employs permutation models and focusses more on ordinary language²⁶. Here, I focus on the so-called permutation argument.

The model-theoretic conception of a theory is one whereby there is a language L with a given interpretation function I, which maps the expressions of L onto a world of objects and properties. Putnam places some methodological constraints on what can fix the intended interpretation. First, there are what Putnam terms ‘theoretical constraints’. These include the standard axioms of set theory, as well as principles and theories from other branches of science. Second, there are the ‘operational constraints’. These are the various empirical observations and measurements that we make in the course of scientific investigation. There is a dual constraint on the assignment of extensions to the subsentential components of the sentences of the language. First, this must be constrained by facts about usage, understood in a suitably naturalistic way. I

²⁴ This claim is contingent on how finely you individuate the conclusions. When I say similar in spirit, therefore, I mean Putnam uses both arguments to attack metaphysical realism.

²⁵ See for example, Putnam (1980).

²⁶ See for example, Putnam (1981) and (1977).

take it that by naturalistically acceptable Putnam intends to acknowledge his Quinean heritage and thus this means taking a somewhat behaviouristic view of the data for semantic theorising where the relevant evidence has to be public. Second, assignments must be constrained by getting the coarse-grained truth conditions of whole sentences correct. The data which must constrain our semantic theorising are the semantic values of whole sentences, and the success condition for a theory of reference is to fit this data set. Importantly, if there are multiple theories which equally fit the data, then it is indeterminate which theory is correct.

The argument then shows that for every model of any theory, it is possible to find a permutation function K , such that each item in L is interpreted “in violently different ways, each of them compatibly with the requirement that the truth value of each sentence in each possible world be the one specified”. (Putnam 1981, 33) The mechanism functions by having the “crazy assignments of reference to names can be ‘cancelled out’ by a compensating assignment of extensions to predicates, so that, overall, the truth value of sentences is unaffected.” (Williams 2007, 369)

What we find is therefore worrying. There are multiple interpretations, all of which meet the requisite constraints. Different interpretations of a sentence can produce the same truth value at every possible world, but assign different extensions to the subsentential expressions. That is: in both the intended and permuted interpretations, ‘cat’ is true of radically different things, but true in exactly the same circumstances nonetheless. Given how we defined the constraints on our metasemantic theorising, neither interpretation can be said to be any better than any other, and it is therefore indeterminate which is correct.

This is deeply in conflict with the commitments of Putnam’s realist. Given the realist’s commitment to uniqueness, we would have hoped that there would be a single true interpretation function which takes us from terms of the theory to the world. However, we find that reference is *radically* indeterminate. This is not the innocuous sort of indeterminacy we considered at the start. It is the claim that even with all the constraints on reference assignment in place, and

even for precise languages like mathematical language, reference is still indeterminate.

These considerations put pressure on another supposed commitment of realism: the non-epistemic nature of truth. Putnam's target realist needs it to be the case that there is an independence of ideal theory and true theory. The realist wants to claim that the intended interpretation is the one that matches the parts of the theory to the right parts of the world. That this is the true theory. So, for the realist there may be ideal theories, which meet all the operational and theoretical constraints, which nonetheless are false. Putnam takes his argument to have shown, however, that such independence cannot be plausibly maintained. He claims:

“Now, such a model satisfies all operational constraints, since it agrees with OP. It satisfies those theoretical constraints we would impose in the ideal limit of inquiry. So, once again, it looks as if any such model is "intended"-for what else could single out a model as "intended" than this? But if this is what it is to be an "intended model", T1 must be true-true in all intended models! The metaphysical realist's claim that even the ideal theory T1 might be false "in reality" seems to collapse into unintelligibility.” (Putnam 1980, 473-474)

For any consistent theory, that theory has a model, and Putnam's argument shows us that if there is one truthmaking model, then there are infinitely many permuted variants. The realist wants to be in a position to say that of all these interpretations, one of them is getting it right and the others are getting it wrong. For Putnam, the intended interpretation is the one that satisfies all the operational and theoretical constraints. Thus we identify truth as idealisation of rational acceptability. If the realist wants to maintain that there is more to theory being true than being intended in Putnam's sense, then they need to provide some means of discriminating between the interpretations. Putnam's contention is that the sort of relation or constraint required by the realist to maintain their distinction between ideal theory and true theory requires commitment to some kind of magical relation or constraint. The realist is therefore not in a position to say anything principled about why one model is getting it right and the others are getting it wrong.

IV. Truth and Metaphysics

In the last two sections I presented two different cases for the claim that truth is epistemically constrained, and we should therefore reject metaphysical realism. In this final section I will evaluate this position. The two main arguments in Dummett that get us epistemically constrained truth are the manifestation and the acquisition arguments, and these, taken together with the claim that the issue at stake between the realist and anti-realist is bivalence gets you the rejection of realism. Putnam, on the other hand, establishes both of these commitments with the model theoretic argument and the claim that the realist requires particular conceptions of truth and reference. It is important to separate out two different kinds of arguments at work. First, there are the arguments for the claim that truth is epistemically constrained. Second, there are arguments which aim to show that this claim about truth is incompatible with the commitments of realism, and so realism should be rejected.

All of these arguments and claims have proved to be contentious. Against both Dummett and Putnam's characterisation of realism, philosophers have argued that they both significantly misrepresent the requirements of metaphysical realism. Specifically, in ascribing to the realist a particular conception of truth, Dummett and Putnam have incorrectly identified the locus of the metaphysical realist's commitments. I will return to consider this line of objection later.

With respect to Dummett's arguments for epistemically constrained truth, the acquisition argument has come under attack most. The main thrust of this attack is from Wright as follows:

“[T]he realist seems to have a very simple answer. Given that the understanding of statements in general is to be viewed as consisting in possession of a concept of their truth-conditions, acquiring a concept of an evidence-transcendent state of affairs is simply a matter of

acquiring an understanding of a statement for which that state of affairs would constitute the truth-condition. And such an understanding is acquired, like the understanding of any previously unheard sentence in the language, by understanding the constituent words and the significance of their mode of combination.” (Wright 1993, 16)

Another line of attack contends an almost Moorean point: *we just do* understand sentences with evidence transcendent truth conditions. In a similar spirit, some philosophers have taken it to be legitimate to merely reject the arguments as the almost uniform acceptance of classical logic in mathematics and science provides a good inference to the best explanation argument for the realist conception of truth. Of course, Dummettian responses are available.

With respect to Putnam, some philosophers have argued that perhaps a more sophisticated theory of reference could show that we were wrong. Putnam’s response to this is that a more sophisticated theory of reference cannot save our intuitive notion of reference as whatever additional constraint we bring in to save some realist notion of reference, could be likewise subjected to a permutation.²⁷ That is: the new constraint itself needs interpretation and each model will interpret it in different ways. The point of Putnam’s argument is that no first order theory can, by itself, determine its own objects up to the point of isomorphism. Putnam states the point thus: “The problem as to how the whole representation, including the empirical theory of knowledge that is a part of it, can determinately refer is not a problem that can be solved by developing a more and better empirical theory.” (Putnam 1980, 477) There are some philosophers who hold that the arguments for radical indeterminacy theses must not be in good standing as they are self-defeating. For example, Scott Soames claims that if the conclusions of the arguments were true, we couldn’t even state them. (Soames 1998, 213) There are other philosophers who have argued that the arguments are not in good standing as they are too stringent in what they allow to be taken as the constraints for providing a theory of reference. Lewis, for example, argues that we need not be worried as

²⁷ For more in-depth discussion of this aspect of Putnam’s argument see Button (2013).

some meanings are more eligible than others and so the type of indeterminacy in question does not arise. (Lewis 1984, 227)

There is of course a great deal more to be said about these responses. For my present purposes, however, I want to precede by separating two different kinds of claims. First, that truth is epistemically constrained and we must therefore reject metaphysical realism. Second, that *if* truth is epistemically constrained *then* we must therefore reject metaphysical realism. If it can be shown that truth is sometimes recognition transcendent, then, given that both Dummett and Putnam have a monistic conception of truth, the first claim turns out to be false. However, the conditional can still hold and, whilst this won't require that we actually need to reject realism, it can establish something significant about the relationship between metaphysics and truth. So in evaluating Dummett and Putnam's positions I will not primarily be concerned with engaging with their arguments for epistemically constrained truth. What I will rather do is (1) highlight the significant point of departure in thinking about truth from the other philosophers and frameworks considered, and (2) articulate what I take to be the metaphysical consequences of holding that truth is epistemically constrained.

iv. Truth and Revisionism

The aim of this section is to map out some foundational issues about how to think about the permissible range of analyses of the concept of truth. The way of thinking about truth that I've detailed in this chapter departs significantly from the positions previously considered. Dummett and Putnam both argue for revisionary conceptions of truth. I take this divergence to be in large part a result of the different constraints placed on the concept, and this raises interesting questions about how we should determine the conceptual role for truth.

This can perhaps best be seen by offering a comparison with the way that Lewis thinks about the concept. As I argued in Chapter One, Lewis's project promises robust conceptions of both meaning and metaphysics, with a central

and integrated role for truth in both these projects. In Lewis's system, truth is both indispensable to semantics, and anchors his theory of meaning, and it supervenes on being, and as such is connected to his metaphysical theorising. Much of what Lewis says about truth is explicitly or implicitly inconsistent with the revisionary approach advocated by Dummett and Putnam. I suggest that what would motivate Lewis's rejection of the sorts of arguments and positions that Dummett and Putnam forward are two interconnected commitments. Lewis takes the conceptual role of truth to be constrained jointly by the role that it plays in the theory of meaning and in metaphysics. Thus, one commitment that precludes the sort of revisionary conception offered by Dummett and Putnam is that, whatever the concept of truth is, it needs to be such that it can service Lewis's metaphysics. A second and related commitment is that Lewis takes it as a constraint on the range of acceptable analyses of any concept that departure from our folk theoretical understanding of that concept counts as a cost. Lewis's Mooreanism, then, provides a significant obstacle to accepting any view of truth according to which there is no gap between what is warrantably assertible and what is true.

So, against Dummett's position I suggest that Lewis would assert that when speakers utter sentences like 'every even integer greater than 2 can be expressed as the sum of two primes', we should take them to be uttering a sentence which is either true or false, and this is independent of our ability to verify it. Dummett forwards a sceptical position about classical semantics. Once we take the manifestation and acquisition arguments seriously, we see that there could not be evidence transcendent truth, which means we need to reject bivalence. Lewis's commitment to Mooreanism is going to preclude his acceptance of Dummett's arguments. Lewis takes it as a datum that classical logic is central to the practice of natural science. As such, anything that makes us doubt it must have gone wrong somewhere.²⁸

When discussing Putnam's attack on realism Lewis states the following:

"If I am looking in the right place for a saving constraint, then realism needs realism. That is: the realism that recognises a nontrivial

²⁸ This is of course not to say that classical logic is in fact *indispensable* to the practice of science.

enterprise of discovering the truth about the world needs the traditional realism that recognises objective sameness and difference, joints in the world, discriminatory classifications not of our own making.” (Lewis 1984, 67)

Lewis offers a central place for truth in the theory of meaning. He also thinks that the theory of meaning needs to take seriously the linguistic practices of speakers; that our theory of meaning/content should make central a theory of communication which is governed by conventions of truthfulness and trust. However, in addition to this, truth also plays an important role in the metaphysics. Lewis’s project of semantic analysis is a prerequisite for metaphysical theorising. Semantic analysis furnishes us with truth conditions for which we require worldly satisfiers; the analysis therefore, informs both metaphysical and empirical inquiry. Lewis thus offers us an integrated conception of truth; it is operationally constrained by the role that it must play in the theory of meaning and in the metaphysics. So, unlike both Putnam and Dummett, who think that the proper understanding of the concept of truth is going to be primarily determined by the role that truth plays in the theory of meaning and interpretation, Lewis takes it to be a legitimate constraint on the concept that it is such that it can also play a role in a common-sense realist metaphysic.

There are two kinds of issues here. All parties agree that the demands placed on the concept of truth by the theory of meaning constrain the acceptable analyses of the concept. So, the first question is, what kind of constraints can be legitimately placed on the concept in the theory of meaning? A second further question is, can legitimate constraints be placed on the concept from outside of the theory of meaning? For example, is it legitimate for our prior metaphysical commitment to constrain the acceptable analyses of the concept? I take it that these considerations raise important and interesting questions about how we should begin theorising about truth. It is not my aim here to determinately answer these questions, but rather to elucidate the theoretical options.

One possible reaction to these considerations is to claim that these issues are so fundamental that we aren’t in the right kind of territory for giving

arguments which will convince your interlocutor one way or the other. I'm imagining a kind of dialectical impasse, where any supposed premise of an argument for your preferred view is the negation of a conclusion from a disputant's argument. In such a case perhaps it is permissible to merely state what your commitments are and then start to theory build from there. Whatever the fundamental truth about truth is, it will be beyond ken. So, it counts as no cost against Dummett's and Putnam's revisionary conceptions of truth that they depart from alethic orthodoxy.

A second option may be to claim that we have some kind of prior, theory neutral, grasp of the concept of truth. The job is then to articulate the features of this concept. Any proposed analysis of the concept of truth will succeed inasmuch as it conforms with this theory neutral concept. Inasmuch as the analysis departs from the theory-neutral concept, this counts against the analysis. So, Dummett's and Putnam's revisionism would only count against their views if it would be shown that their view falls short of this theory neutral conception. Such a view of course invites many more questions. Is there really such a theory neutral understanding available of the concept of truth which we can articulate? Perhaps a good candidate for such an understanding would be Wright's platitudes? Again, further questions require answering. Which platitudes? Perhaps we just have the core ones, like disquotation, but to make the conception theory neutral we must reject many of the others, like bivalence?

A final option would be to claim that we shouldn't be in the business of trying to find out the fundamental truth about truth. It seems plausible that there just isn't a consistent concept which satisfies all the roles we may want to assign to it. You might indeed take this to be the moral of the Liar Paradox: there just isn't a concept which can satisfy all the demands we want to place upon it. So, assuming you think that there is interesting important work for truth, then this will require being a little ameliorative about quite what this concept is. Such view takes truth as a theoretical concept, and then a candidate analysis of the concept is judged relative to the theoretical work it can do. By way of comparison, think about Lewis on meaning. Such an approach contends that we aren't trying to articulate the fundamental truth about what meaning is.

Rather, what we are trying to do is to demarcate a philosophically useful concept which we can put to work. And then the important question for my purposes is: given this analysis of the concept, what consequences, if any, does this have for metaphysics?

There is of course much more to say about all of these options. However, I do not attempt to settle these fundamental issues in this chapter. I will proceed by assuming that an epistemically constrained conception of truth is a legitimate way to think about the concept, and the arguments in favour of it are in good standing. In the next section I turn to address the metaphysical consequences of thinking about truth in the way proposed by Dummett and Putnam.

v. Evaluating the Conditional

The aim of this section is to evaluate the supposed connection between truth and realism. Dummett and Putnam each think that it is a constitutive thesis of realism that truth is not epistemically constrained. So, when they show that truth is epistemically constrained, we have to reject metaphysical realism. One way of resisting this argument is to show that truth is epistemically unconstrained. Another route is to undermine the connection between truth and realism. This latter approach is the one I investigate here. In particular, I want to defend the following conditional: if truth is epistemically constrained, then this has unavoidable metametaphysical consequences. Defending this conditional requires some significant revision of what Dummett and Putnam each take to be the realist/anti-realist consequences of their respective arguments.

My plan is as follows. First, I state my assumptions about truth. I assume that (1) truth is epistemically constrained, (2) the theory of meaning is truth-conditional, and (3) a monistic conception of truth. Second, I argue that we can articulate a conception of metaphysical realism which is entirely independent of any considerations to do with the nature of truth. Finally, I argue that despite the availability of a conception of realism which is not in

part constituted by an alethic commitment, the epistemically constrained nature of truth nonetheless has metaphysical significance.

My starting assumptions are as follows. First, I assume that truth is epistemically constrained. I will assume that the arguments given by Dummett and by Putnam for that conclusion are sound. As I've already indicated, those arguments have come under attack from multiple directions. Nonetheless, to adequately make the point I want to make about the relationship between truth and metaphysics I will proceed under the assumption that truth is in fact epistemically constrained. Related to this assumption, I want to shelve another issue which, while relevant, is orthogonal to my present argument. If we assume the soundness of Putnam's argument for epistemically constrained truth as I present it here, we are accepting an argument for radical semantic indeterminacy. As I argue elsewhere (Fletcher 2016, 49-51) I think that Putnam's argument for radical semantic indeterminacy has serious implications for metaphysical realism. However, assessing these implications of semantic indeterminacy for metaphysical realism will not be my concern in this section.

Second, I assume that truth has an ineliminable role to play in the theory of meaning. A way out of the argument that I forward below is to claim that we needn't employ a truth conditional theory of meaning. If we have available a theory of meaning which does not employ the concept of truth, then we can reject the thought that the correctness of a theory must consist in it being ideally warrantably assertible. Throughout this thesis I have assumed that truth has some indispensable role to play in the theory of meaning, so I don't take my assumption at this point to be controversial.

Finally, in addition to these assumptions, I further assume that there is just one truth concept. There is a live debate about the viability of alethic pluralism. If I am successful in showing that if truth is epistemically constrained then this creates a certain kind of challenge to realism, then perhaps one way that this could be resisted is to argue that whilst there is a truth concept which is epistemically constrained, there are others which are not and that one of these

robust correspondence concepts is the one being employed by the realist.²⁹ One way of putting this is to say that we can locally accept the conditional. There are some domains where the relevant truth concept must be construed as epistemically constrained, however, there are other where it is not. For example, perhaps the relevant notion of truth pertaining to moral statements is epistemically constrained, and so this has consequences for the objectivity of moral discourse. By contrast, perhaps the relevant notion of truth pertaining to the statements of best science is robust correspondence, and so we can legitimately talk about the statements of scientific theories being such that their truth values are evidence transcendent. Whilst I think that this debate is an interesting one, I will proceed by talking about *the* concept of truth, and the metaphysical consequences of *this* concept.

Now I turn to articulate a conception of metaphysical realism which is entirely independent of any considerations to do with the nature of truth. One way to neutralise Dummett's and Putnam's attacks on realism is to show that truth has nothing to do with realism. This line is adopted by Michael Devitt. Devitt asks this question: "what does truth have to do with realism?" (Devitt 1983, 292) His answer to the question of what truth has to do with realism is "nothing at all". (Devitt 1983, 292) Devitt captures the notion of realism he is interested in, and according to him the one metaphysicians should be interested in, in terms of two commitments. These commitments run along two dimensions. First: existence. Realism commits us to a view about what exists. If I am a realist about some entity, or class of entities, I am committed to those things existing. The second commitment is trickier: "words that frequently occur in attempts to capture the second are 'independent', 'external' and 'objective'. The entities must be independent of the mental; they must be external to the mind; they must exist objectively in that they exist whatever anyone's opinions." (Devitt 1983, 292)

The doctrine of realism according to Devitt should be construed thus:
"common sense, and scientific, physical entities objectively exist independent

²⁹ I do not consider alethic property pluralism because it is far from clear how an epistemically constrained concept could pick out an epistemically unconstrained property.

of the mental.” (Devitt 1983, 292) Construed as such, I can see the temptation to claim that truth does not have anything to do with realism. It makes no mention of semantic notions at all and it does not seem clear, at least *prima facie*, how this characterisation might implicitly rest on the sort of alethic commitments required to get Dummett’s and Putnam’s attack going. Devitt claims that the type of metaphysical significance that Dummett and Putnam take their arguments to have is only possible if you endorse a language first approach to metaphysics.

Perhaps we can view Devitt’s criticism as a plea not to conflate metaphysics and the philosophy of language. It is only by conflating the two that one might be inclined to think that considerations from the philosophy of language could impact upon the realist status of the world. Of course, language is a tool for theorising. We do theorise *using* language, but we must not let this confuse us. The *content* of our theorising has nothing to do with language, truth or any kind of semantic concerns, and so we should not be alarmed by Dummett’s and Putnam’s putative challenge. Metaphysics, realist metaphysics, is concerned with the world, not language.

Additionally, perhaps we can see some immediate appeal to Devitt’s position. There is intuitive appeal; what we can and cannot say, what we can and cannot theorise about, should not affect what actually is the case. That there is a world, and that that world exists mind-independently, should not be affected by any concerns pertaining to the epistemic limitations of creatures like us. In misrepresenting the commitments of realism Dummett and Putnam get away with far more than they should. Whilst we may still have cause for concern about the significance of Dummett and Putnam style arguments in the case of metasemantics, this does not force us to any conclusions about the viability of metaphysical realism. Realism, so construed, has minimal commitments (i.e. an existence thesis and an independence thesis), and neither of these commitments are affected by the semantic concerns forwarded by Dummett and Putnam.

An obvious concern one could have about Devitt’s proposal is that we don’t have an alethically neutral way of explaining in what the independence thesis consists. Dummett in particular is worried that we don’t have a non-

metaphorical way of explicating what independence consists in, and this is what motivates him to claim that it should be understood in terms of bivalence. What is going to be really important for Devitt, therefore, in maintaining the plausibility of this alethically neutral conception of realism is that the independence commitment can be spelled out without talk of truth. I think that, at the very least, there are intelligible conceptions of independence which do not require taking a view about the nature of truth. So, I proceed assuming that Devitt has offered an alternative way of conceiving of the commitments of realism.

Both Putnam and Dummett hold the view that the philosophy of language (or specifically the theory of meaning or interpretation) underpins debates about realism. Now, what if you think, like Devitt, that metaphysical realism has nothing to do with truth? Does the view of truth discussed in this chapter nonetheless show us something about the limits of metaphysics? I close this chapter by arguing that it is infelicitous of Devitt to say that truth has nothing to do with realism. The reason for this is because the requirements of realism are threefold, and Devitt only secures one such commitment.

What does it take for a theory to be a realist? The metaphysical doctrine of realism according to Devitt requires that there is an independently existing reality. However, if I claim to be a realist metaphysician I take it that there are two further methodological presuppositions which are necessary for me to engage in theorising in the first place. First, I must think that I have some kind of access to the world. This could be seen as an epistemological presupposition. That is: we can have knowledge of an independently existing part of reality. Second, I must think that I am capable of representing the world in my theories. This could be viewed as a metasemantic presupposition. Whatever I take the *metaphysical* characterisation of realism to be, it seems plausible that I must hold these two presuppositions as given if my metaphysical theories are to do what any realist would want them to do.

The challenge I forward here accepts Devitt characterisation of metaphysical realism, and, as such, holds that Dummett's and Putnam's conception of truth cannot have *metaphysical* significance. However, I argue that it does have great significance for the prospect of realist *theorising*. It should be clear, therefore,

that the picture I propose is not one that endorses any kind of anti-realism; I am not making any claims about what exists or the mind-independent/dependent nature of what exists.

How does epistemically constrained truth create a problem for realist theorising? If truth is epistemically constrained, then what is true cannot outstrip what is warrantably assertible.³⁰ So, there is a certain sense in which the semantic presupposition and the epistemic presupposition collapse into each other, in that we have to establish what is knowable in order to establish what is representable. This matters for the following reason. Realists very often want to shelve the epistemology of metaphysics. However, if truth is epistemically constrained, then they have to provide an adequate epistemology in order to provide a theory at all. In particular, what the realist cannot do is construct a theory and claim that it is true or false but we cannot know which because what it is for a statement to have truth conditions is for that statement to be knowable.

An objection to my line of thinking here would be that no metaphysician is going to claim that we have no way of knowing whether a metaphysical theory is true. First, it is worth noting that if I am a Williamson-style epistemicist about vagueness³¹, and I think that there is some truth about the cut off point for baldness but that we can never know it, then this provides at least one example of the sort of theory which would be off limits if truth were epistemically constrained. However, I suspect that many metaphysicians would claim that the answers to metaphysical questions are knowable. In particular, I suspect that a metaphysical realist would claim that someone who thought that epistemically constrained truth could make trouble for realism in the way I've suggested, has an inadequate conception of what can count as evidence. According to this view, the claims of serious metaphysics are knowable because the serious metaphysician can employ theoretical virtues which could provide warrant for asserting one metaphysical thesis over another.

³⁰ We might prefer to follow Wright and put this in terms of superassertibility (see Wright 1992), but for my current purposes I'm going to continue talking in terms of warranted assertibility.

³¹ See Williamson (1994).

It is beyond the scope of this thesis to settle the issue of the evidential status of theoretical virtues. Of course, if it turns out that the sort of cost-benefit analyses, employing theoretical virtues, which are frequently given as a reason to adopt one metaphysical theory over another, do not constitute an acceptable source of evidence, then those ‘theories’ turn out to not be sensibly stateable. So, whilst the limits of our knowledge are not the limits of the world, the limits of our knowledge to provide a limit to metaphysical theory. Even if we think that adopting the sort of ‘metaphysics-first’ approach advocated by Devitt is desirable, an inescapable question for realist metaphysics remains: can an adequate epistemology of realist metaphysics be given? In conclusion, therefore, if truth is epistemically constrained, then this has unavoidable metametaphysical consequences.

V. Conclusion

The aim of this chapter was to evaluate the supposed connection between truth and realism. Dummett and Putnam both think that it is a constitutive thesis of realism that truth is not epistemically constrained. So, when they show that truth is epistemically constrained, we have to reject metaphysical realism. I investigated a possible way of resisting their conclusion which proceeds by arguing that there is no such connection between truth and realism. I concluded that whilst severing the link between metaphysical realism and truth can save some intuitive conception of realist metaphysics, this move is unsuccessful in recovering a realist conception of metaphysical theories.

Conclusion

The aim of this thesis was twofold. First, I aimed to ascertain whether there is a consistent and explanatorily potent concept of truth which can play a role in the theory of meaning and in metaphysics. Second, I aimed to determine, having thus established a role for truth in metaphysics, the conception of metaphysics with which this particular role for truth furnishes us. I aimed to provide a theoretical road map which detailed the commitments, choice points, and assumptions pertinent to different approaches I considered. I indicated throughout places where more work needs to be done or where argumentative moves require further justification. In this conclusion, therefore, I want to recap, pointing to the challenges that each position faces in sustaining an acceptable role for truth in metaphysics, and to describe what I take the significant lessons to be.

In Chapter One I argued that in adopting Lewis's project of philosophical analysis we could have robust conceptions of both meaning and metaphysics, with a central and integrated role for truth in both of these projects. I presented Lewis's project of philosophical analysis as the means by which both the semantic and metaphysical projects should be prosecuted and how they connect. Lewis's approach offers a promising project for those metaphysicians who want to bridge the gap between the ways in which we ordinarily think about the world and what best science and metaphysics tell us about what the world is really like. I argued that the primary purpose of philosophical analysis is to give the meanings of the concepts under investigation, and that this enterprise is a prerequisite for successful metaphysical theorising. The role of truth in metaphysics is exhausted by delivering the semantic analyses that are required to identify the subject-matter on which metaphysics goes to work and by allowing us to forge a connection to being. Thus to ensure a robust metaphysics, independent metaphysical methods must be utilised. I argued that whether such an approach is acceptable is dependent upon providing justification for this independent methodology.

A recurring theme in my discussion of Lewis, Sider, and the realist response to Dummett and Putnam was the epistemological status of theoretical virtues. All of those approaches hope to utilise a method for metaphysical theory construction which borrows the criteria for selecting between competing theories from science. If the trajectory of the claim is that, because these are legitimate methods in science then they can be legitimate methods in metaphysics, then there are two places where pressure can be applied. First, we might wonder whether science and metaphysics are sufficiently similar that utilising the same tools is appropriate. Second, in order for metaphysics to inherit such methods, and consequently justification, it needs to be the case that the theoretical virtues really play the role suggested in science. There is a serious and contentious question in the philosophy of science about the evidential status of theoretical virtues, and so even if metaphysics and science can be shown to be sufficiently similar, there is still a question about whether these methods reliably yield true theories.

In Chapter Two I argued against Dyke's metaphysical conception of truth. Dyke holds that if we want to adequately sustain a role for truth in metaphysics, we need to cut this notion off from concerns that primarily pertain to our systems of representation. The approach is therefore promising if you are worried that issues in semantics and metasemantics will get in the way of serious metaphysical work. Nonetheless, I argued that if we propose to make use of the concept of truth in a metaphysical framework that distances itself from the role assigned to truth in the theory of meaning, then we must have a story to tell about how that concept can play the role assigned to it. I further urged that we must be cautious that our desire to do serious metaphysics doesn't rule out the possibility of robust metaphysical projects that are anti-realist, and that we must not allow a move towards a 'metaphysics-first' method to rob us of the ability to actually prosecute metaphysical theorising. Finally, I argued that when thinking about metaphysical method we must be careful to distinguish between claims of epistemological priority versus metaphysical priority.

Part of the discussion in this chapter highlighted a very important separation of methodological claims from metaphysical ones. I think that when thinking

about the role of truth in metaphysics this distinction cannot be emphasised enough. Thinking that metaphysics has reason to care about truth and meaning, as these notions apply to theories of natural language, is apt to generate criticism from those metaphysicians who think that metaphysics is about the world, not language or its properties. Thus, it is crucial when thinking about these matters that a distinction is made between the centrality of semantic considerations to metaphysical *methods* versus semantic terms having their place (be it central or otherwise) in metaphysical *conclusions*. The type of role for truth that has been the preoccupation of this thesis is how truth is relevant to the method of metaphysics, how that impacts upon the status of metaphysics.

In Chapter Three I argued that Sider presents a robustly realist conception of metaphysics, where the epistemological status of the enterprise of metaphysics should be seen as on par with that of natural science. I argued that truth plays a significant role in Sider's proposed method for metaphysics, first, via the role that truth plays in connecting the nonfundamental to the fundamental, and second, via the idea of truth as a required feature of metaphysical theory. In a similar spirit to Dyke, Sider thinks that in order to sustain the role of truth in metaphysics, we need to have available a notion which does not carry the same semantic and metasemantic baggage that it does in natural language. Sider's project presents an attractive method for those metaphysicians who want a completely general approach to metaphysical inquiry. Nonetheless, I concluded that Sider insufficiently motivates his assigned role for truth, and thus a central part of Sider's project is left worryingly opaque.

The discussion of Sider's rejection of empirical semantics projects highlighted some very important issues in thinking about the role of truth in metaphysics. I argued that those who adopt something like an empirical semantics project hold that semantic concepts only earn their keep if they play some role in empirical theories and that this requires that there have to be discernible connections between the application of these concepts and observable facts. So, in the case of truth, truth earns its keep in an empirical theory of communication. Adding the concept of truth to our conceptual repertoire makes an explanatory difference when we are trying to give empirical theories

of understanding, communication, belief, and interpretation. In particular, we only get a respectable notion of truth if it is bound to observable facts about assent. Truth is that thing that sentences are held to have when they are asserted or believed. If the only way we can motivate a role for truth in theorising simpliciter is through its role in the theory of meaning, then any role that it may play in metaphysics is going to need to be accordingly constrained by this primary function.

In a similar spirit to the conclusion of the previous chapter, then, the important take away from this chapter is that if we propose to make use of the concept of truth in a metaphysical framework then we must have a story to tell about what legitimises this use. Recall I said in the Introduction that I was taking it as a constraint on the acceptable deployment of the concept of truth that it cannot make it opaque how truth can play the role assigned to it. Now, it would be uncharitable to claim that Dyke or Sider do not offer truth an explanatory role. Both of them take truth to play a very important role in metaphysics. But, crucially, it is far from obvious how the explanatory role that they assign to truth makes an *empirical* difference to successful theories. I suggest that more work needs to be done in thinking about the legitimacy of deploying the concept in a way which is autonomous from the explanatory role it plays in theorising about the observable facts of human communication. This is not to say that the challenge cannot be met, or indeed that a satisfactory story must follow analogous lines to empirical semantics projects, but I suggest that whether Sider or Dyke's approaches to thinking about truth in metaphysics are acceptable depends on whether some adequate story can be told.

In Chapter Four I argued that Davidson assigns an important role to truth in metaphysics. Davidson advocates for an approach to ontological inquiry according to which we derive ontological commitments when the truth conditional analysis of the true sentences of the language requires quantifiers and domains of quantification. The concept of truth that plays this robust role in metaphysics is one which is shaped with the demands of an empirical semantics project in mind. Importantly, the nature of Davidson's empirical semantics project is such that we cannot ensure subsentential semantic

determinacy. Thus, given the method for metaphysics that Davidson forwards, we cannot talk in *absolute* terms about the ontological commitments of our best theory.

I argue that Davidson's project illuminates how the status of metaphysical inquiry can be transformed depending on claims made about semantic concepts. If you think both that a truth conditional analysis of language is a prerequisite for generating metaphysical commitments, and that the interpretation of that language is indeterminate, then these commitments together have a serious impact upon the type of metaphysical project in which you can engage. For approaches to metaphysics which take the study of language to be in some sense authoritative, Davidson's project is instructive in demonstrating how assumptions about the nature of truth and the theory of meaning can transform the limit or status of metaphysics. I therefore caution that we must be mindful of how metaphysical assumptions can get smuggled in to the method of metaphysics. I take both the discussion of Davidson and of Dummett and Putnam to usefully demonstrate how seemingly benign assumptions about truth and reference can have an impact upon the status of metaphysical conclusions. In particular, in this context, the claims that content is determinate or that truth is bivalent are not metaphysically innocent.

In Chapter Five I argued that if the metaontological deflationist approach to ontological inquiry is to be workable, it needs to be furnished with an adequate account of truth. I showed how truth plays a significant role in the metaontological deflationist framework: truth is that feature which we take the sentences from which we derive our ontological commitments to have. I further articulated how the approach would be attractive to those philosophers who want to demystify the method of metaphysics. Nonetheless, I claimed that as it stands, an adequate account of truth has not been sufficiently articulated. Given the significant theoretical role that truth plays, this should be worrying. On the one hand, making good on the epistemological and methodological virtues of the approach is held hostage to the development of such an account. On the other hand, and more importantly for my current purposes, for the metaontological deflationist's conception of the role of truth in metaphysics to be compelling, they need to be able to provide some kind of story about how it

is that truth can play the role assigned to it. I suggested that an adequate notion could be articulated, and I gestured towards an approach I believe to be a fruitful in attempting to do just that. Of course, the acceptability of this way of thinking about truth is going to depend on whether it is an acceptable way of thinking about the role of truth in the theory of meaning.

In Chapter Six I argued if truth is epistemically constrained, then this has unavoidable metametaphysical consequences. Dummett and Putnam each think that it is a constitutive thesis of realism that truth is epistemically unconstrained. So, when they show that truth is epistemically constrained, we have to reject metaphysical realism. I investigated a possible way of resisting their conclusion which proceeds by arguing that there is no such connection between truth and realism. I concluded that whilst severing the link between metaphysical realism and truth can save some intuitive conception of realist metaphysics, this move is unsuccessful in recovering a realistic conception of metaphysical theories. Thus, an important take away from this this chapter was similar to that of Chapter Two: we must ensure that a ‘metaphysics-first’ method does not rob us of the ability to actually prosecute metaphysical theorising.

A significant consideration that was highlighted in this chapter, which came out of the discussion of how Dummett’s and Putnam’s projects differed from the others I have considered, was about the legitimate demands that can be placed on the concept of truth. I claimed that it is possible that there might not be a consistent concept which can satisfy all of the roles that we may want to assign to it. So, assuming you think that there is interesting and important work for truth, this may require being a little ameliorative about the concept. What the limits of amelioration are with respect to truth is an open question, and one that requires further study. Whatever we say about these matters, however, I hope to have demonstrated how we must be sensitive to the complex sets of demands that we place on the concept of truth. In particular, we must be mindful of how the role assigned to truth in the method of metaphysics can limit or impact upon the status of metaphysical inquiry.

Bibliography

- Armstrong, David (1980) 'Against Ostrich Nominalism: A Reply to Michael Devitt', in *Pacific Philosophical Quarterly*, 61.
- Barnes, Elizabeth & Williams, J. R. G. (2011) 'A Theory of Metaphysical Indeterminacy' in Karen Bennett & Dean W. Zimmerman (eds.), *Oxford Studies in Metaphysics* volume 6. Oxford University Press.
- Barnes, Elizabeth (2014) 'Going Beyond the Fundamental: Feminism in Contemporary Metaphysics', in *Proceedings of the Aristotelian Society*, 114.
- Berkovski, Sandy (2011) 'Lewis' Reduction of Modality', in *Acta Analytica* 26 (2).
- Button, Tim (2011) 'The Metamathematics of Putnam's Model-Theoretic Arguments', in *Erkenntnis* 74:3.
- , (2016) 'Deflationary Metaphysics and Ordinary Language', in *Synthese*.
- Cameron, Ross P. (2010), 'How to Have a Radically Minimal Ontology', in *Philosophical Studies* 151.
- Carnap, Rudolf (1950) 'Empiricism, Semantics, and Ontology', in *Revue Internationale de Philosophie*.
- Davidson, Donald (1965) 'Theories of Meaning and Learnable Languages', in *Inquiries into Truth and Interpretation*, 2nd Edition (2001). Oxford University Press.
- , (1967) 'Truth and Meaning', in *Inquiries into Truth and Interpretation*, 2nd Edition (2001). Oxford University Press.
- , (1969) 'True to the Facts', in *Inquiries into Truth and Interpretation*, 2nd Edition (2001). Oxford University Press.
- , (1970) 'Semantics for Natural Languages', in *Inquiries into Truth and Interpretation*, 2nd Edition (2001). Oxford University Press.

- , (1973a) 'In Defence of Convention T', in *Inquiries into Truth and Interpretation*, 2nd Edition (2001). Oxford University Press.
- , (1973b) 'Radical Interpretation', in *Inquiries into Truth and Interpretation*, 2nd Edition (2001). Oxford University Press.
- , (1974) 'Appendix to Essay 10. Replies to Lewis and Quine', in *Inquiries into Truth and Interpretation*, 2nd Edition (2001). Oxford University Press.
- , (1977a) 'The Method of Truth in Metaphysics', in *Inquiries into Truth and Interpretation*, 2nd Edition (2001). Oxford University Press.
- , (1977b) 'Reality Without Reference', in *Inquiries into Truth and Interpretation*, 2nd Edition (2001). Oxford University Press.
- , (1979) 'Moods and Performance', in *Inquiries into Truth and Interpretation*, 2nd Edition (2001). Oxford University Press.
- , (1984) 'Communication and Convention', in *Inquiries into Truth and Interpretation*, 2nd Edition (2001). Oxford University Press.
- , (2001) 'Introduction', in *Inquiries into Truth and Interpretation*, 2nd Edition (2001). Oxford University Press.
- Devitt, Michael (1983) 'Realism and the Renegade Putnam: A critical study of Meaning and the Moral Sciences', in *Noûs* 17 (2).
- Divers, John (2017) 'How Skeptical Is Quine's "Modal Skepticism"?', in *The Monist*, Vol. 100, No. 2.
- Dorr, Cian (2008) 'There are no Abstract Objects', in Theodore Sider, John Hawthorne & Dean W. Zimmerman (eds.), *Contemporary Debates in Metaphysics*. Blackwell.
- Dummett, Michael (1978) *Truth and Other Enigmas*. Harvard University Press.
- , (1991) *The Logical Basis of Metaphysics*. Harvard University Press.
- , (1993) *The Seas of Language*. Oxford University Press.
- Dyke, Heather (2008) *Metaphysics and the Representational Fallacy*. Routledge.
- Fletcher, Jade (2016) 'Arguing About Realism: Adjudicating the Putnam-Devitt Dispute', in *European Journal of Analytic Philosophy*, Vol. 12, No. 2.

- Fine, Kit (1994) 'Essence and Modality', in *Philosophical Perspectives* 8.
- , (2001) 'The Question of Realism', in *Philosophers' Imprint* 1 (1).
- Frege, Gottlob (1980) *The Foundations of Arithmetic*, J. L. Austin (trans.).
Northwestern University Press.
- Hale, Bob (2013) *Necessary Beings: An Essay on Ontology, Modality, and the Relations Between Them*. Oxford University Press.
- Hale, Bob & Wright, Crispin (2009) 'The Metaontology of Abstraction', in David John Chalmers, David Manley & Ryan Wasserman (eds.), *Metametaphysics: New Essays on the Foundations of Ontology*. Oxford University Press.
- Heil, John (2003) *From an Ontological Point of View*. Oxford University Press.
- Jago, Mark (2018) *What Truth Is*. Oxford University Press.
- Keefe, R. (2000) *Theories of Vagueness*. Cambridge University Press.
- Künne, Wolfgang (2003) *Conceptions of Truth*. Oxford University Press.
- Ladyman, James & Ross, Don (2007) *Everything Must Go: Metaphysical Naturalized*. Oxford University Press.
- Lewis, David (1968) 'Counterpart Theory and Quantified Modal Logic', in *Philosophical Papers Volume 1* (1983) Oxford University Press.
- , (1969) *Convention*. Blackwell (2002 paperback edition)
- , (1970a) 'Anselm and Actuality', in *Philosophical Papers Volume 1* (1983) Oxford University Press.
- , (1970b) 'General Semantics', in *Philosophical Papers Volume 1* (1983) Oxford University Press.
- , (1970c) 'How to Define Theoretical Terms', in *Philosophical Papers Volume 1* (1983) Oxford University Press.
- , (1971) 'Counterparts of Persons and Their Bodies', in *Philosophical Papers Volume 1* (1983) Oxford University Press.
- , (1972) 'Psychophysical and Theoretical Identifications', in *Papers in Metaphysics and Epistemology* (1999) Cambridge University Press.

- , (1973a) 'Causation', in *Philosophical Papers Volume 2* (1986) Oxford University Press.
- , (1973b) *Counterfactuals*. Wiley-Blackwell (2001 paperback edition)
- , (1973c) 'Counterfactuals and Comparative Possibility', in *Philosophical Papers Volume 2* (1986) Oxford University Press.
- , (1974a) 'Radical Interpretation', in *Philosophical Papers Volume 1* (1983) Oxford University Press.
- , (1974b) 'Tensions', in *Philosophical Papers Volume 1* (1983) Oxford University Press.
- , (1975) 'Languages and Language', in *Philosophical Papers Volume 1* (1983) Oxford University Press.
- , (1976a) 'Probabilities of Conditionals and Conditional Probabilities', in *Philosophical Papers Volume 2* (1986) Oxford University Press.
- , (1976b) 'Survival and Identity', in *Philosophical Papers Volume 1* (1983) Oxford University Press.
- , (1979a) 'Attitudes De Dicto and De Se', in *Philosophical Papers Volume 1* (1983) Oxford University Press.
- , (1979b) 'Counterfactuals Dependence and Time's Arrow', in *Philosophical Papers Volume 2* (1986) Oxford University Press.
- , (1980) 'A Subjectivist's Guide to Objective Chance', in *Philosophical Papers Volume 2* (1986) Oxford University Press.
- , (1981) 'What Puzzling Does Not Believe', in *Papers in Metaphysics and Epistemology* (1999) Cambridge University Press.
- , (1983a) 'Extrinsic Properties', in *Papers in Metaphysics and Epistemology* (1999) Cambridge University Press.
- , (1983b) 'New Work for a Theory of Universals', in *Papers in Metaphysics and Epistemology* (1999) Cambridge University Press.
- , (1983c) 'Postscript to Anselm and Actuality', in *Philosophical Papers Volume 1* (1983) Oxford University Press.

- , (1983d) 'Postscript to Attitudes De Dicto and De Se', in *Philosophical Papers Volume 1* (1983) Oxford University Press.
- , (1984) 'Putnam's Paradox', in *Papers in Metaphysics and Epistemology* (1999) Cambridge University Press.
- , (1986a) 'Against Structural Universals', in *Papers in Metaphysics and Epistemology* (1999) Cambridge University Press.
- , (1986b) *On the Plurality of Worlds*. Wiley-Blackwell (2001 paperback edition)
- , (1986c) 'Postscript to A Subjectivist's Guide to Objective Chance', in *Philosophical Papers Volume 2* (1986) Oxford University Press.
- , (1986d) 'Postscript to Causation', in *Philosophical Papers Volume 2* (1986) Oxford University Press.
- , (1986e) 'Postscript to Counterfactuals Dependence and Time's Arrow', in *Philosophical Papers Volume 2* (1986) Oxford University Press.
- , (1988) 'Rearrangement of Particles: Reply to Lowe', in *Papers in Metaphysics and Epistemology* (1999) Cambridge University Press.
- , (1992) 'Armstrong on Combinatorial Possibility', in *Papers in Metaphysics and Epistemology* (1999) Cambridge University Press.
- , (1994a) 'Humean Supervenience Debugged', in *Papers in Metaphysics and Epistemology* (1999) Cambridge University Press.
- , (1994b) 'Reduction of Mind', in *Papers in Metaphysics and Epistemology* (1999) Cambridge University Press.
- , (1996) 'Elusive Knowledge', in *Papers in Metaphysics and Epistemology* (1999) Cambridge University Press.
- , (1997a) 'Finkish Dispositions', in *Papers in Metaphysics and Epistemology* (1999) Cambridge University Press.
- , (1997b) 'Naming the Colours', in *Papers in Metaphysics and Epistemology* (1999) Cambridge University Press.

- Livingston, Paul M. (2005) 'Functionalism and Logical Analysis', in David Woodruff Smith & Amie L. Thomasson (eds.), *Phenomenology and Philosophy of Mind*. Oxford: Clarendon Press.
- Maguire, Barry (2013) 'Defending David Lewis's Modal Reduction', in *Philosophical Studies*, 166 (1).
- Melia, Joseph (1995) 'On What There's Not', in *Analysis*, 55 (4).
- Menzies, Peter & Price, Huw (2009) 'Is Semantics in the Plan?' in David Braddon-Mitchell & Robert Nola (eds.), *Conceptual Analysis and Philosophical Naturalism*. MIT Press.
- Merricks, Trenton (2001) 'Varieties of Vagueness', in *Philosophy and Phenomenological Research*, 63.
- Price, Huw (2009) 'Metaphysics After Carnap: The Ghost Who Walks?', in David John Chalmers, David Manley & Ryan Wasserman (eds.), *Metametaphysics: New Essays on the Foundations of Ontology*. Oxford University Press.
- Putnam, Hilary (1977) 'Realism and Reason', in *Proceedings and Addresses of the American Philosophical Association* 50.
- , (1978) *Meaning and the Moral Science*. Routledge and Kegan Paul Ltd.
- , (1980) 'Models and Reality', in *Journal of Symbolic Logic*, 45 (3).
- , (1981). *Reason, Truth, and History*. Cambridge University Press.
- Quine, W. V. (1969) *Ontological Relativity and Other Essays*. Columbia University Press.
- , (1960) *Word & Object*. The MIT Press.
- Ritchie, Jack (2008) 'Structuralism Realism and Davidson', in *Synthese*, 162 (1).
- Reichenbach, Hans (1938) 'On Probability and Induction', *Philosophy of Science* 5, no. 1: 21–45. Reprinted in S. Sarkar, ed. *Logic, Probability and Induction*, Garland, New York, 1996.
- Sider, Theodore (2009), 'Ontological Realism', in David Chalmers, David Manley, and Ryan Wasserman (eds.), *Metametaphysics* (Oxford: Oxford University Press).

- , (2011) *Writing the Book of the World*. Oxford University Press.
- Soames, Scott (1998) 'Skepticism about meaning, indeterminacy, normativity, and the rule-following paradox' in *Canadian Journal of Philosophy* 23.
- Tarski, Alfred (1936) 'The Concept of Truth in Formalized Languages, in *Logic, Semantics, Metamathematics* (1983). Oxford University Press.
- Thomasson, Amie (2014) 'Deflationism in Semantics and Metaphysics', in *Metasemantics: New Essays on the Foundations of Meaning*, (eds.) Burgess, Alexis & Sherman, Brett. Oxford University Press.
- , (2015) *Ontology Made Easy*. Oxford University Press.
- Thompson, Naomi (2016) 'Grounding and Metaphysical Explanation', in *Proceedings of the Aristotelian Society* 116 (3).
- van Fraassen, Bas (1980) *The Scientific Image*. Oxford University Press.
- Williams, J. R. G. (2007) 'Eligibility and Inscrutability', in *Philosophical Review* 116 (3).
- Williamson, Timothy (1994) *Vagueness*. Routledge.
- , (2014) 'How Did We Get From There to Here? The Transformation of Analytic Philosophy', in *Belgrade Philosophical Annual*, 27.
- Wright, Crispin (1992) *Truth and Objectivity*. Harvard University Press.
- , (1993) *Realism, Meaning, and Truth*. Oxford: Blackwell.