

|

A meta-ontological criticism of Eli Hirsch's semanticist attack on
physical object ontology

Carl Simon Murphy

This thesis is presented to Cardiff University in partial fulfilment of the
requirements of the degree of Doctor of Philosophy

30 September 2016

DECLARATION

This work has not been submitted in substance for any other degree or award at this or any other university or place of learning, nor is being submitted concurrently in candidature for any degree or other award.

Signed (candidate) Date
.....

STATEMENT 1

This thesis is being submitted in partial fulfillment of the requirements for the degree of(insert MCh, MD, MPhil, PhD etc, as appropriate)

Signed (candidate) Date
.....

STATEMENT 2

This thesis is the result of my own independent work/investigation, except where otherwise stated, and the thesis has not been edited by a third party beyond what is permitted by Cardiff University's Policy on the Use of Third Party Editors by Research Degree Students. Other sources are acknowledged by explicit references. The views expressed are my own.

Signed (candidate) Date
.....

STATEMENT 3

I hereby give consent for my thesis, if accepted, to be available online in the University's Open Access repository and for inter-library loan, and for the title and summary to be made available to outside organisations.

Signed (candidate) Date
.....

STATEMENT 4: PREVIOUSLY APPROVED BAR ON ACCESS

I hereby give consent for my thesis, if accepted, to be available online in the University's Open Access repository and for inter-library loans **after expiry of a bar on access previously approved by the Academic Standards & Quality Committee.**

Signed (candidate) Date
.....

Abstract

Physical object ontology is a sub-branch of ontology which is primarily concerned with three interrelated issues. These are the composition, material constitution, and manner of persistence for physical objects. Philosophers who take up positions on such issues often disagree over what objects they think the world contains – for example a mereological nihilist will argue that there are no composite objects, and thus would say that the ordinary objects that appear to be all around us do not actually exist. Such disputes are thought to be substantive and depend for their truth on what the world itself is actually like.

Against this Eli Hirsch develops a meta-ontological argument which states that the debates in physical object ontology are merely verbal, and that what is going in these debates is that each side is simply speaking an alternate language in which their claims come out trivially true and the claims of their opponent come out trivially false. Thus there is no actual disagreement over the facts. This position of Hirsch's I call semanticism.

The purpose of this thesis is to articulate Hirsch's position, demonstrating its Carnapian roots, but also showing how Hirsch, by making several key commitments, intends his position to be distinctive from a thoroughgoing Carnapianism and its potentially unattractive commitments to anti-realism and/or verificationism. However, in this thesis I develop a number of problems for Hirsch's position, showing that his modified version of Carnapianism is untenable, and that he is forced between giving up his central contention or retreating into a more thoroughgoing Carnapianism.

Acknowledgements

The patience and support of many individuals has made this thesis possible. I would in particular like to thank my supervisor Professor Nicholas Shackel for his knowledge, guidance and support. I would also like to thank Rhian Rattray for having the patience of a saint and for being generally fantastic.

Contents

Chapter one: Introduction	1
1: Plan of the thesis	1
2: A brief reminder of physical object ontology	3
3: A note on abbreviations	6
Chapter two: Eli Hirsch, semanticism and quantifier variance	7
1: Introduction	7
2.1: Introducing quantifier variance	8
2.2: How does quantifier variance work?	10
2.3: Some points of contention	13
3: Some consequences arising from quantifier variance	15
3.1: Quantifier variance and a deflationary theory of reference	15
3.2: Translation and compositional truth-theoretic semantics	20
4: Top-down semantics and neo-Fregeanism in Hirsch	24
4.1: The question of how sentences get their character and the notion of a 'linguistic decision'	26
4.2: Bringing together Hirsch's thoughts	28
4.3: More points of contention	31
5: Quantifier variance and the shallow approach to ontology	32
6: Why accept quantifier variance?	34
6.1: The plausibility argument part I	35
6.2: The plausibility argument part II	40
6.3: Some comments on the plausibility argument	42
7: Characterising Hirsch's position	43
7.1: Summary of chapter two and general outline of the argument of the thesis	45
Chapter three: Bennett and the analytic problem	50
1: Introduction	50
2: A preliminary worry: What are the sufficient conditions for a merely verbal debate?	51
3: 'LP' has to be analytic in the CSO-language	55
4: 'LP' can't be analytic in the CSO-language	59
5: Saving Hirsch from linguistic idealism	62
6: The CSO speaker does not consider 'LP*' to be analytic – the 'analytic problem'	66
6.1: The verificationist retreat from the analytic problem	71
7: Summary of chapter three	74
Chapter four: A potential solution to the analytic problem	76
1: Introduction	76
2: The first solution: Do not commit to 'AS' being analytic or synthetic	78
3: The second solution: The dispute over 'AS' itself is merely verbal	83
4: Problems with explaining the truth-values of attitude sentences	86
5: Tractarian-analyticity	89
6: Does the commitment to T-an pose a problem for Hirsch?	93
7: Summary of chapter four	100
Chapter five: Hawthorne's criticisms of Hirsch	102
1: Introduction	102
2: The intension-centric approach to semantics	102
3: The first intensional worry	103
4: The second intensional worry	105
4.1: The statue vs. the lump ontologist and the nihilist vs. the anti-nihilist	106

<i>4.2: Gunkism vs. atomism</i>	<i>111</i>
<i>4.3: Endurantism vs. Perdurantism</i>	<i>112</i>
<i>5: Summary of the intensional worries</i>	<i>114</i>
<i>6: Hyperintensional issues</i>	<i>115</i>
<i>7: Summary of chapter five</i>	<i>122</i>
Chapter six: The problem of hyperintensionality for Hirsch	124
<i>1: Introduction</i>	<i>124</i>
<i>2: Fine grained content in Hirsch</i>	<i>125</i>
<i>3: Two accounts of hyperintensionality</i>	<i>129</i>
<i>4: What kind of hyperintensionality does Hirsch require?</i>	<i>131</i>
<i>5: The CSO speaker does not regard the hyperintensional differences to be merely representational</i>	<i>135</i>
<i>6: Potential solutions for Hirsch</i>	<i>137</i>
<i>7: Summary of chapter six</i>	<i>140</i>
Chapter seven: The Sider/Hirsch debate	142
<i>1: Introduction</i>	<i>142</i>
<i>2: Meaning is determined by use plus eligibility</i>	<i>144</i>
<i>2.1: External constraints on the meanings of quantifier expressions</i>	<i>150</i>
<i>3: Hirsch's reply to Sider</i>	<i>153</i>
<i>4: Assessing the debate thus far</i>	<i>157</i>
<i>5: Sider's reply to Hirsch</i>	<i>159</i>
<i>6: Hirsch's second reply to Sider</i>	<i>164</i>
<i>7: Sider's second reply</i>	<i>170</i>
<i>8: Assessing the debate</i>	<i>175</i>
Chapter eight: A final problem for Hirsch	180
<i>1: Introduction</i>	<i>180</i>
<i>2: Language relates primarily to the world at the level of sentences</i>	<i>181</i>
<i>3: The nature of facts</i>	<i>184</i>
<i>4: The argument to monism</i>	<i>186</i>
<i>5: A way of avoiding monism: Tractarian factualism</i>	<i>189</i>
<i>6: A final option: Epistemicism</i>	<i>193</i>
<i>7: Summary of the argument in this chapter</i>	<i>194</i>
<i>8: General conclusion</i>	<i>196</i>
Bibliography	203

Chapter one: Introduction

1: Plan of the thesis

For some time metaphysicians have been happy to engage in substantive disputes about what concrete objects the world contains – a sub-branch of ontological dispute called physical object ontology. However, there has been an increasing scepticism about the substantive nature of these debates, and meta-ontology has seen a resurgence of neo-Carnapian approaches which argue that far from being substantive debates about what objects the world actually contains, the different positions are merely to be understood as different but equally correct and true ways of describing the same reality – i.e. the general point is that there is no actual disagreement between each side over the facts. This is despite what the disputants themselves might actually think. And so just as the logical positivists attempted to use language to show up large tracts of philosophy as consisting of nothing more than nonsense and pseudo-disputes, the neo-Carnapians attempt to show certain areas of ontological debate as reducing to nothing more than linguistic decisions.

The most prominent proponent of neo-Carnapianism is Eli Hirsch. The essential claim that Hirsch makes is that each side in a purported ontological dispute is speaking a different language in which their asserted sentences come out true, and the asserted sentences of their opponent come out false. This then allows each side to charitably interpret the other as speaking the truth in their own language. This then determines that neither side is actually in any disagreement over the (language independent) facts themselves.

The goal of this thesis is to critically examine Hirsch's position. Specifically I will argue that whilst Hirsch's position is certainly neo-Carnapian, Hirsch seeks to distance himself from Carnap by making a number of explicit commitments (I will discuss just what these are in chapter two). This has the effect of making Hirsch Carnapian in spirit but not to the letter – this is *neo*-Carnapianism after all. However, I will demonstrate in the thesis that the additional commitments that Hirsch makes in an effort to distinguish his position from a thoroughgoing Carnapianism leads him in to significant trouble. Ultimately, I will demonstrate that Hirsch cannot maintain his core claim that the disputes in physical object ontology are merely verbal whilst maintaining these additional commitments. Thus Hirsch must be seen to either give up his central claim, or give up these additional

commitments and retreat toward a more traditional, thoroughgoing Carnapianism.

In what remains of this chapter I will briefly recount the main positions in physical object ontology. In chapter two I will then give an in-depth account of Hirsch's position. This chapter will be largely exposition, but I also offer some key insights about Hirsch's position that he himself has not really shone sufficient light on – such as the fact that he appears to be committed to a deflationary theory of reference. I also draw together the various claims Hirsch makes in the three articles I discuss into an overall cohesive position, showing how claims that Hirsch makes in his earlier articles can be incorporated into claims he makes in his later articles, which is something he does not do himself. Then I look at the precise ways in which Hirsch seeks to differentiate his position from thoroughgoing Carnapianism.

In chapter three, I discuss an article by Bennett and bring to light a problem that I have called the *analytic problem*. The analytic problem, though not called this by Bennett, was originally developed by her. In this chapter I add some important arguments to Bennett's position which she has overlooked, and show how some of her points can be met by Hirsch if we pay particular attention to his commitments in the philosophy of language which I analysed in chapter two. Though it is not possible to go into the nature of the analytic problem here, this chapter establishes it as a serious issue for Hirsch. It explores some tentative solutions on behalf of Hirsch, all of which I argue are unsatisfactory as they amount either to an abandonment of Hirsch's central claim or signify a retreat into a thoroughgoing Carnapianism.

In chapter four I develop on behalf of a Hirsch a sustained solution to the analytic problem. It is again not possible to go into specifics here but the proposed solution, though initially attractive, can only work if Hirsch gives up on some of his key commitments, showing that the analytic problem can only be avoided by retreating toward thoroughgoing Carnapianism.

In chapter five I move on to discuss a number of criticisms of Hirsch as presented by Hawthorne. However, I demonstrate that two main arguments that Hawthorne presents are not effective criticisms of Hirsch. I then point out that Hawthorne raises an interesting point regarding the use of hyperintensional operators in metaphysics, which are purportedly used to track structural features of reality. According to Hawthorne, such operators cannot play that intended role owing to

Hirsch's views on language. However I argue that this point in itself does not constitute an argument against Hirsch – it is rather an undeveloped observation.

Chapter five thus serves as a springboard to chapter six where I develop a new problem for Hirsch's position based on the notion hyperintensionality, which I have called the H-problem. The H-problem is shown to be problematic for Hirsch in the same way that the analytic problem was. I then show in this chapter how the H-problem can only be solved by retreating further into thoroughgoing Carnapianism.

In chapter seven I discuss the lengthy debate in the literature between Sider and Hirsch, since Sider has positioned himself as Hirsch's most prominent critic. A summary of this dispute is valuable in its own right, but I ultimately demonstrate that Sider's criticisms fail to adequately account for Hirsch's commitments in the philosophy of language, and thus his criticisms fall short.

This then serves as a springboard for the final chapter, in which I argue that if we take into proper consideration Hirsch's commitments in the philosophy of language, we can still have room for ontological debate. I then point out a very obvious block that Hirsch has against this move. However, I then demonstrate that ultimately Hirsch is forced into a position where he must choose between two problematic and substantive commitments, or he can retreat into thoroughgoing Carnapianism, or he can modify his position to incorporate epistemological claims. Either way, the only way Hirsch can avoid the problems presented in this final chapter is by modifying his position beyond what he originally intended it to be.

It will be seen by the end of the thesis that Hirsch's particular version of neo-Carnapianism is not sustainable. This means that we can either have substantive disputes in ontology, or we can have thoroughgoing Carnapianism, but we cannot have Hirsch's modified version of Carnapianism. This still leaves substantive approaches to ontology in need of defending of course, but it should be seen by the end of this thesis that Hirsch's particular position, as originally presented, is untenable.

2: A brief reminder of physical object ontology

Physical object ontology incorporates a number of interrelated positions regarding the composition and manner of persistence of physical objects. Revisionary

positions regarding matters of composition mean accepting that there will be far more or far fewer physical objects than we ordinarily assume.

The main positions one can take up with regards to composition questions are the following:

Mereological nihilism: According to the nihilist, there are no composite objects – the world only contains physical simples (objects which have no proper parts). The ordinary objects we see around us are, on this view, merely arrangements of simples – there is nothing ‘over and above’ the simples in their various arrangements. Thus on this view the world contains far fewer objects than we hitherto recognise in ordinary thought and language. An example of an argument in favour of nihilist would be Merrick’s *causal over determination argument*. This argument basically states a composite object, if it exists, would cause only the same thing as its parts acting in concert. If composites exist then, their effects would be over determined. This makes composite objects causally redundant and for this reason Merrick’s thinks they need not play a role in our ontology (Merricks, 2001: 58)¹.

Universalism: According to the universalist, for any two objects x and y , there is some further object z which has x and y as proper parts. Thus composition never fails to occur. The result of this view is that the world contains far more objects than we hitherto have recognised, including strange scattered objects such as the thing composed of Clinton’s nose and the Eiffel Tower. In his 2001a, Sider argues that we should accept universalism. He argues that if it is not the case that for any xs , the xs always compose an object, then either there will be arbitrary cut-off points in composition, or composition would be vague. Since both of these results are counter-intuitive, the only option left is that composition always occurs (Sider 2001a, 122). One might wonder why Sider doesn’t accept nihilism as an acceptable response to the worry of vagueness and arbitrary cut-offs in composition. It could be argued that universalism, for all the counter-intuitive objects it posits, at least saves the objects of common sense from elimination whereas nihilism does not.

Common sense: The common sense ontologist aims to defend our ordinary understanding of the kinds of objects the world contains. Lowe is an example of a proponent of a common sense approach to ontology. His method is to make use of

¹ Merricks argues that humans luckily escape this argument, but the argument itself can be taken as a good example of an argument in favour of eliminating composite objects.

other traditional metaphysical categories such as kinds which supply composition and persistence conditions for objects (Lowe, 2006: 50). In his 2003 he also offer an argument to show that Merrick's over determination argument can actually be seen to let in far more objects than Merrick's admits by showing that composite objects instantiate properties that cannot be understood simply as the cumulative result of the properties of its parts (Lowe, 2003: 708).

Another example of a defender of common sense would be Markosian, who argues that facts about composition are brute. He argues that if we are forced to trade on unintuitive propositions, the proposition that composition is brute is not as bad as accepting a position like nihilism or universalism, thus by elimination, brute composition is the best position (Markosian, 1998).

The second important issue in physical object ontology is the manner in which physical objects persist through time. It is commonly held that objects can undergo qualitative change over time whilst persisting through this change. Ordinarily it is supposed that physical objects are just the kind of things that can persist through certain kinds of change. However, some philosophers feel this is in need of deeper metaphysical explanation.

Perdurantism: Perdurantists accept an ontology of *temporal parts* whereby ordinary physical objects, as well as being composed of spatial parts, are also composed of temporal parts. On this view of the world, an object is not 'wholly present' at any point of time, but is *spread out* through time as well as space. All that exists at each point of time is a temporal part of the object. The object exists in virtue of being composed out of these temporal parts.

Four-dimensionalism: Perdurantists often also accept the thesis of unrestricted composition – in this thesis I will refer to such a position as four-dimensionalism. Sider is a proponent of four-dimensionalism arguing that issues about vagueness and arbitrary cuts off in composition arise diachronically, and thus we should accept that for any object or set of objects at t_1 and for any other object or set of objects at t_2 , there will be a cross-time fusion of those objects (i.e. the fusion will have those objects, at different times, as temporal parts) (Sider, 2001a: 133-4).

Endurantism: Against the perdurantist/four-dimensionalist picture of the world we have the common-sense approach which argues that objects do not have temporal parts, and that any objects that exists is wholly present at each time that it exists.

Lowe argues that temporal parts are an ontologically extravagant flight of fancy – we do not need to posit them to explain how objects change over time (Lowe 1988: 77). That objects can undergo a change of parts and persist through other kinds of qualitative change is to be explained by reference to the *kind* that the object instantiates – each kind supplies *persistence conditions* for objects that instantiate that kind (Lowe 1988: 76).

3: A note on abbreviations

In this thesis I will often shorten the expression ‘common sense ontology’ to CSO, and will refer to someone who adopts this position as a CSO. As I will explain fully in chapter two, Hirsch claims that each ontologist can be regarded as speaking the truth in his own language. Thus the CSO claim that ‘there is a table here’ will come out true in the CSO-language, for example. Hirsch imagines then that each position has associated with it a language in which the assertions of that position comes out true. For the nihilist, I abbreviate this as the N-language frequently throughout the thesis. Anyone who is taken to speak this language I will refer to as an N-speaker, and this applies to the other languages also. Finally, I abbreviate the four-dimensionalist language to the 4d-language. At other times I will quote directly from Hirsch who sometimes might refer to 4d-English, for example. In those cases the context will make it obvious which language Hirsch is taking about.

Chapter two: Eli Hirsch, semanticism and quantifier variance

1: Introduction

Against the view that debates in physical object ontology are substantive, Eli Hirsch has become the leading proponent of a view labelled by some (such as Hawthorne 2009) as semanticism, or by others as semantic deflationism (Manley 2009). Hirsch takes his position to be something of a continuation of Putnam's thesis of internal realism (Hirsch 2002), but also describes his position as 'roughly Carnapian' (Hirsch 2009, 231), owing something to Carnap's internal/external distinction (Hirsch 2005, 67). Whilst Hirsch considers his position to have this much in common with Carnap, he considers his position to be distinct from Carnap's in several important respects. The ways in which he takes his position to differ from Carnap's will be discussed in section 7 of this chapter, and it will be seen that the tensions that arise in Hirsch position, to be explored in later chapters, are a direct result of his attempt to construe his position as what I see as a modified 'light' version of Carnapianism. It is ultimately the aim of this thesis to demonstrate that this 'light' variety of Carnapianism is untenable. I will have more to say on this in section 7 of this chapter, but for now I will simply explicate Hirsch's position as he presents it. The core argument he presents is that there is nothing substantively at stake in the debates in physical object ontology, and that what is actually going on is that the disputants are in effect describing the same facts in reality using different languages, unbeknownst to themselves. As such, they are not really disagreeing about how the world is itself, but are in actuality engaged in a pseudo-dispute which should resolve itself once it has been pointed out to them that they are speaking different languages in which their asserted sentences come out true. Thus, for example, whereas a nihilist will say that 'there is nothing in region R but simples arranged table-wise', the common-sense ontologist will disagree and say 'there is a table in region R '. According to Hirsch, the nihilist sentence will come out true in the nihilist language, whereas the common-sense sentence will come out true in the common-sense language, and therefore neither side has spoken falsely and has *made no factual mistake about the world itself*. Thus whilst the ontologist will say that whether or not the sentence 'there is a table' is true depends on how the world is, what Hirsch says is that this sentence will be true or false depending on which language we elect to speak. The consequence of this is that ontologists engaged in disputes in physical object ontology should not be seen as being engaged in a substantive dispute about what

objects the world itself contains, but instead should be seen as doing nothing more than engaging in a pseudo-dispute. This has distinct consequences for where and how ontological discussion should proceed, which I will discuss in section 5 of this chapter.

In this chapter I will first explain Hirsch's notion of quantifier variance, which is of central importance to claim that ontologists are speaking different languages. After this, I will show that a commitment to quantifier variance involves two major commitments – a commitment to a deflationary theory of a reference, and a commitment to a repudiation of a referential correspondence theory of truth. These will then be explained in the context of Hirsch's later stated endorsement of a 'top-down' approach to semantics, as opposed to a 'bottom-up' approach. Thus the first half of this chapter will be devoted to explaining the technical machinery which Hirsch uses in order to advance his argument. Explaining *what* quantifier variance is one thing, arguing for *why* we should accept it as a sound notion is another, and thus in the second half of the chapter I will explain how Hirsch uses quantifier variance to advocate a so-called 'shallow' approach to questions of ontology, and will also then discuss his argument as to why we should accept the notion of quantifier variance.

If Hirsch's argument is right, then it presents a serious challenge to physical object ontology as it is currently practiced. Throughout this chapter at relevant points I will draw attention to areas of tension within Hirsch's position which others have picked up on, and these will then be used as springboards to develop criticisms of Hirsch's particular brand of semanticism in later chapters.

2.1: Introducing quantifier variance

In this section I'm going to introduce the notion of quantifier variance and explain how Hirsch intends to use it to show that it the debates that constitute physical object ontology can be regarded as verbal. A more detailed account of just how quantifier variance is supposed to work is given in section 2.2.

The central claim of quantifier variance is that natural language quantifier expressions such as 'there exists' can have a variety of different meanings, and whether we choose to adopt one meaning over another will affect the truth values of sentences in which the quantifier expression is used. The result of this is that by reinterpreting the meaning of the quantifier 'there exists', we can interpret the

sentence ‘there is a table’, for example, in a way that makes it true or in a way that makes it false. As Hirsch says “the quantificational apparatus in our language and thought - such expressions such as ‘thing’, ‘object’, ‘something’, ‘(there) exists’ – has a certain variability or plasticity” (Hirsch, 2002: 51). He goes on to add that ‘there is no necessity to use these expressions in one way rather than various other ways, for the world can be correctly described using a variety of concepts of “the existence of something” (Hirsch, 2002: 51). This is important in establishing that debates in physical object ontology are merely verbal. Consider the following from Hirsch:

Suppose we are evaluating the truth of the sentence, ‘there exists something that is composed of Clinton’s nose and the Eiffel Tower’. Mereologists will accept this sentence, whereas anti-mereologists will reject it...quantifier variance implies that the expression ‘there exists something’ can be interpreted in a way that makes the sentence true or in a way that makes it false. Since both interpretations are available to us, we have a choice of operating with a concept of ‘the existence of something’ that satisfies the mereologist or operating with a different concept that satisfies the anti-mereologist. (Hirsch, 2002: 51)

If we imagine then a universalist and a common-sense ontologist arguing over whether there is an object that is composed of Clinton’s nose and the Eiffel Tower, Hirsch will say that this is merely a pseudo-dispute, since the answer to that question will have a different answer according to which concept of existence we choose to use. This creates a problem for ontologists who think that the question is substantive and that the answer depends on how the world is in itself. This kind of ontologist will argue that whether or not the sentence ‘there exists something that is composed of Clinton’s nose and the Eiffel Tower’ is true or false will depend on how things stand in objective reality, whereas Hirsch will argue that the sentence will be true or false depending on which concept of existence we choose to use. Since that choice is up to us, if two people are disagreeing over the truth value of the sentence, the disagreement arises not because of a disagreement over how the world is itself, but arises merely because both sides are unknowingly using different concepts of existence. Thus Hirsch suggests that what is going on in ontological disputes of this kind is that each side is speaking a particular language,

L , where their asserted sentences come out true and the sentences asserted by their opponent come out false. However, they can understand their opponent as speaking a different language, L^* , such that in L^* the opponent's asserted sentences do indeed come out true, owing to the fact that L^* operates with a different concept of the existence of a thing. As Hirsch says, 'each side can plausibly interpret the other side as speaking a language in which the latter's asserted sentences are true' (Hirsch, 2009: 231). Hirsch thinks then that the disputes between the nihilist, common-sense ontologist and universalist and the dispute between the endurantist and perdurantist can all be understood as pseudo-disputes, since each side can be seen as speaking the truth in their own language (Hirsch 2005: 68-89). It is quite clear then that what makes these different languages possible is the notion of quantifier variance – the idea that the quantificational apparatus of a language has a variety of interpretations which can make certain sentences come out true or false depending on what interpretation we are working with.

2.2: How does quantifier variance work?

Thus far I have introduced the notion of quantifier variance and how it is used by Hirsch to show that debates in physical object ontology are merely verbal. The question now is this: just what does it mean to say that we can use the quantifier in a variety of different ways? And, following on from that, how can we imagine a different way of using the expression 'there exists' in a way such that it still is a quantifier expression – for we can easily suppose that 'cats love to chew bones' is true if we just reinterpret 'cat' to mean *dog*, but then of course we are no longer talking about *cats*. Thus any such re-imagining of the meaning of a quantifier must be such that whatever comes out the other end is still legitimately a quantifier expression.

First then, consider that Hirsch thinks we can only understand the meaning of the quantifier by understanding the role it plays in determining the truth-conditions of the sentences in which it used, in the same way that we understand the meaning of logical connectives. Thus he says that 'quantifier variance is not a matter of substituting one "definition" for another; it's a matter of substituting one range of truth-conditions for another' (Hirsch, 2002: 54). He then explains that 'In general, we explain the meaning of a logical constant by describing the role it plays in determining the truth-conditions of sentences. Thus we explain the meaning of

“and” by saying that sentences of the form “p and q” are true if and only if both the sentence “p” and the sentence “q” are true’ (Hirsch, 2002: 54). The idea here is that someone who understands that if ‘p’ is true and ‘q’ is true then ‘p and q’ is true understands the meaning of the expression ‘and’. Hirsch argues that this is the same as how we should understand the meaning of quantificational expressions – we look to what role they play in determining the truth-conditions of sentences. We can explain a shift or difference in meaning of the quantifier expressions by explaining the different roles the quantifier can have in determining the truth-conditions of sentences. Let’s look at an example from Hirsch:

In my use of “there exists something” a sentence of the form “There exists something composed of the F-thing and the G-thing” counts as true only if “the F-thing” and the “G-thing” refer to things that are connected (united) in some special ways. Let me call my use of “there exists something” the A-use. I now imagine a different use of “there exists something” – I will call this the M-use – in which a sentence of that form counts as true so long as the “F-thing” and the “G-thing” refer to things, no matter how they are connected. (Hirsch 2002: 54)

What Hirsch is saying here is that there are two different truth-conditions associated with the sentence “There exists something composed of the F-thing and the G-thing” depending on whether we adopt the A-use of the quantifier or the M-use. Under the M-use, the truth conditions of the sentence are such that so long as the expressions ‘the F-thing’ and ‘the G-thing’ refer, the sentence is true. Under the A-use then, this is all that is required. However, under the A-use, it appears that in order for the sentence to count as true, the F-thing and the G-thing must be united in some special way. Let us say that *R* designates this relation. Thus on the A-use, the sentence will only be true if ‘the F-thing and the G-thing are *R*-related’ is also true. Hirsch construes this as different ‘semantic rules’ governing sentences of the form “There exists something composed of the F-thing and the G-thing” (Hirsch 2002: 54-5). As such he says that ‘I assume...that this explicit stipulation carries with it a natural way of filling in the truth-conditions for an indefinite variety of other sentences in the M-language’ (Hirsch 2002: 55).

The idea then is that quantifier variance entails that *there are different semantic rules governing sentence forms involving quantificational expressions*. Thus we can speak a language which is governed by one rule as opposed to another rule, and as such in our language a sentence of the form 'There exists an F' may be true whereas in another language, owing to *nothing other than* a difference in semantic rules, 'there exists an F' will be false.

In the example given above then, we can describe the semantic rule governing the M-use of the quantifier as follows:

(R1): For any two expressions, 'x' and 'y', if 'x' and 'y' refer then 'there exists something composed of x and y' is true

Thus we can understand the *M-language* as being a language where R1 is a semantic rule, whereas the *A-language* is a language where R1 is not a semantic rule. Let's imagine two people who are disagreeing over the truth value of the sentence 'there exists something that is composed of Clinton's nose and the Eiffel Tower'. Mr. M thinks this is true, whereas Mr. A thinks it is false. Both, however, agree that the expressions 'Clinton's nose' and the 'Eiffel Tower' refer. Thus we can understand their disagreement as stemming from the fact that Mr. M is speaking the M-language, whereas Mr. A is speaking he A-language. Because R1 is a semantic rule in the M-language, the fact that 'Clinton's nose' and 'the Eiffel Tower' both refer mean that the truth-conditions for the contested sentence are fulfilled, whereas under the A-use the truth-conditions for the contested sentence are not fulfilled (since Clinton's nose and the Eiffel Tower are not R-related). What is striking to notice then is that both Mr. M and Mr. A will agree that Clinton's nose exists and that the Eiffel Tower exists, and they will both also agree that these two are *not* R-related. Thus they appear to be in agreement about a lot of things. What Hirsch is now suggesting is that the point where the apparent disagreement arises can be explained by the different truth conditions associated with the contested sentence, and that once this is pointed out, it will be seen that there is in fact no disagreement between Mr. M and Mr. A about *how the world itself actually is*. This stems from the idea that Mr. A can make intelligible to himself the way that Mr. M speaks by understanding that Mr. M is speaking a language where R1 is a semantic rule. Conversely, Mr. M can make intelligible to himself the way that M. A is speaking by understanding Mr. A is speaking a language where R1 is *not* a semantic rule. This explains what Hirsch means then when he says that 'each side

can plausibly interpret the other side as speaking a language in which the latter's asserted sentences are true' (Hirsch 2009: 231).

The basic idea then is that we can make sense of there being different languages that are characterised by different semantic rules for sentence forms involving quantifier expressions, such that in one language a sentence can be true owing to a particular semantic rule, whereas in another language the same sentence is false owing to different semantic rules. From the perspective of our language, we can make sense of the assertions of a person speaking another language by understanding the different semantic rules that are operative within the language that they are speaking. Once we do this we will see that we are not actually disagreeing about how the world is.

Of course an important issue now remains – we have explained quantifier variance by supposing that there can be different languages such that in each language there are different semantic rules for sentence forms involving quantifier expressions, and that this constitutes a different use, or meaning, of the quantifier in different languages. But as mentioned at the start of this section, to what extent can we say in each case that the expression 'there exists' is still functioning in the different languages as a *quantifier*?

Consider that we can re-interpret the symbol '&' in such a way that if 'A' is true and 'B' is true then 'A & B' is false. We can do this, but then of course we are no longer talking about *conjunction*, but some other logical function. What Hirsch argues is that in each case, in each language the 'the formal logical properties of the expression will not be changed at all (the formal principals of quantificational logic will be unaltered)' (Hirsch 2002: 53). Though he doesn't go into any more detail on this, we can suppose him to mean that the quantifier expressions can be introduced in the standard way. Thus if ' a is F ' is true in L , then 'there exists an x such that x is F ' will be true in L . The idea seems to be then that if this core function of the quantifier is maintained across languages then it is right to say that these expressions still count as quantifier expressions.

2.3: Some points of contention

At this juncture some basic points of contention can be raised. These issues with Hirsch's position form the basis for some of the criticisms levelled against him by Bennett and Sider, which will be expanded upon in chapters three and seven.

The first issue is concerning the understanding of these so-called semantic rules. We have seen, for example, that in an important sense it is a result of a semantic rule in the M-language that the truth of ‘there exists something composed of Clinton’s nose and the Eiffel Tower’ follows necessarily from ‘the expression “Clinton’s nose” refers’ and ‘the expression “the Eiffel Tower” refers’. Thus it appears that the only reason the sentence is true is because *it simply follows from the rules of the language itself*. However, a legitimate question can be asked here – do the speakers of the M-language themselves regard the truth of the sentence as simply following from the rules of their language, or do they think it follows necessarily because it is a *metaphysically necessary truth*, i.e., does the truth of the sentence express something necessary about the world itself, not simply the way we have to describe the world in order to be speaking the M-language? Bennett explores whether such an issue creates a problem for Hirsch’s claim that the debates in physical object ontology are merely verbal. This will be explored in chapter three. Following on from this, I will argue in chapter four that there are potential solutions to the problems raised by Bennett, but that none of them turn out to be satisfactory for Hirsch (see section 7.1 of this chapter for more detail on this point).

Sider also has an issue with Hirsch at this point. Consider the closing paragraph of section 2.2. It appears here as if Hirsch is just assuming we can come up with any set of ‘semantic rules’ involving quantifier expressions in order to then characterise the different languages needed in order for Hirsch’s position to work. However, Sider questions the legitimacy of such a move by considering that there are *external constraints on what quantifier expressions can mean*. If there are external constraints on what quantifiers can mean, it would turn out that the sort of semantic rules Hirsch requires wouldn’t really be legitimate rules for sentence forms involving quantifier expressions, and thus the languages required for Hirsch’s position to work couldn’t exist. If Sider is right about this then Hirsch’s position cannot even get off the ground because quantifier variance is not a legitimate option. This idea forms the starting point for Sider’s criticism which I will discuss in chapter seven.

Now let’s return to discussing some of the consequences of quantifier variance.

3: Some consequences arising from quantifier variance

It has been seen that the core claim of quantifier variance is that we can make sense of there being different languages that are characterised by different semantic rules for sentence forms involving quantifier expressions, such that in one language a sentence can be true owing to a particular semantic rule, whereas in another language the same sentence is false owing to different semantic rules. We have also seen how this can be used to show that the debates in physical object ontology are merely verbal. However, accepting quantifier variance generates two related issues – the first issue is that quantifier variance will appear to generate problematic contradictions involving the notion of reference. I examine this problem in section 3.1. A related issue, explored in 3.2, is that if these different languages are indeed possible, then from the perspective of one language it will not always be possible to provide a compositional truth-theoretic semantics for supposedly true sentences of other languages. This then means we are unable to translate supposedly true sentences of other languages into true sentences of our own. This is problematic if the different languages are taken to be describing (truly) the same facts in reality. I deal with these issues historically in the order that they are discussed by Hirsch. However, in sections 4 – 4.2 I discuss Hirsch’s later commitment to a top-down, neo-Fregean approach to semantics, and show in section 4.2 how the problems about reference, compositional semantics and translation appear to be better dealt with once we understand Hirsch as being committed to a top-down neo-Fregean approach to semantics².

3.1: Quantifier variance and a deflationary theory of reference

As we have seen thus far, Hirsch maintains that a sentence such as ‘there is a table’ can be interpreted in a way that makes it true and in way that makes it false, and this depends only on what semantic rules governing sentence forms containing quantifiers are operative in a language. Thus in the CSO-language, the sentence ‘there is a table’ will be true, where as in the N-language the sentence will be false.

² It is quite possible that Hirsch always had something like a top-down approach to semantics in mind but never thought it needed precise stating or spelling out, or perhaps he later came to realise that the issues facing quantifier variance as highlighted in the following sections required him to adopt a top-down approach to semantics.

This creates a problem if we assume a substantive theory of reference. A commitment to a substantive theory of reference, as I see it, amounts to two central claims:

- 1) The fact that an expression, ' a ', refers, means that there is some relation between the word and an object in reality. Thus there is a substantive reference relation that links the word to language independent reality. The existence of the relation thus entails the existence of both relata – and that includes the object in language independent reality.
- 2) The fact that the sentence 'there is an a ' is true is explained by the fact that the singular expression ' a ' refers – thus the reference of the singular expression determines the truth value of a sentence on which it is used. We can state this as a general principle:

A sentence of the form 'there is an a ' is true iff the expression ' a ' refers.

The fact then that the sentence 'there is a table' can be regarded as true in the CSO-language whilst also being false in the N-language will generate two related contradictions. If 'there is a table' is true, then it implies that the expression 'table' refers, and given (1) above, this means that there is some object in language-independent reality such that there is a relation between the expression 'table' and the object in reality.

However, the sentence 'there is a table' is false in the N-language. Given that the general principle above is a bi-conditional, because 'there is a table' is false, it implies that the expression 'table' does not refer, and by (1) this implies that there is no relation between the word and language-independent reality. Quite simply, thus far it looks like in order for Hirsch's position to work, language-independent reality will have to contain a table and not contain a table at the same time, which is of course absurd. Let's call this the *metaphysical* contradiction because it implies reality itself both contains and does not contain some object.

A related problem is the following. Given that in the N-language the sentence 'there is a table' is false, it is also false in the N-language to say that 'the expression "table" refers'. However, they are supposed to be able to regard 'there is a table' as a true sentence in the CSO-language, as per Hirsch's claim. Thus, given the general principle above, this would seem to imply that 'the expression "table" refers' would have to be true in the N-language, since the sentence 'there is a table' couldn't be

true in *any language* unless the expression ‘table’ referred. Thus the N-speaker seems like he will have to assert that the expression ‘table’ both refers and does not refer at the same time, which again is absurd. Let’s call this the *semantic* contradiction because it implies that in some language, both a sentence and its negation will have to be true at the same time.

A way out of this problem for Hirsch then is to repudiate the substantive theory of reference and replace it with a deflationary one. Whilst Hirsch does not explicitly commit himself to a deflationary theory of reference, I think certain comments he makes certainly support the claim that he appears to assume one. Additionally, there is no way out of the above contradictions unless a deflationary theory of reference is supposed, and thus Hirsch must be committed to a deflationary theory of reference if he wishes for his project to be successful. What I should note at this point however is that a full explanation of how a deflationary theory of reference can block the above contradictions is not possible to give until I have given an account of Hirsch’s commitment to a top-down, neo-Fregean approach to semantics. Thus in section 4.2 I will return to the issues raised in this section and give a complete explanation of why a deflationary theory of reference can block the contradictions stated here.

Firstly then, what is a deflationary theory of reference? A proponent of a deflationary theory of reference will hold that everything there is to know about reference is adequately captured by understanding the following sentence ‘if the sentence “there is an *a*” is true, then the expression ‘*a*’ refers’. We do not require the notion of a substantive relation between an expression and a piece of language independent reality. The deflationist would thus deny the first central claim of a substantive theory of reference. What this means is that notion of reference becomes *nothing but a formal device* that allows us to switch from *using* an expression to mentioning it. Thus accepting a deflationary account of reference involves nothing more than accepting that the schematic sentence

(S) If the sentence “there is an *a*” is true then the expression ‘*a*’ refers

is true in one’s own language. Thus if it is indeed true in *L* that ‘there is *a*’, then it would appear to follow *trivially* that the sentence ‘the expression ‘*a*’ refers’ is also a true sentence of *L* – and this follows simply from the fact that, as an interpreted sentence, ‘there is an *a*’ is *already* true in *L*, and that ‘S’ is also true in *L*. As such, the truth value of the sentence ‘there is an *a*’ isn’t determined or explained by the

truth of the sentence ‘the expression ‘ a ’ refers’ – rather this sentence appears to simply follow from what is already true in L . Thus the notion that the reference of the expression ‘ a ’ somehow is explanatorily prior to the truth of the sentence ‘there is an a ’ is repudiated according to a deflationary theory of reference, and thus the second central principle of a substantive theory of reference is also denied. This ties in with Hirsch’s commitment to neo-Fregeanism, which will be discussed in section 4.

How does this now allow us to block the metaphysical and semantic contradictions that appear to arise? First of all, because the deflationary theory of reference does not commit to there being a substantive relation between an expression and language independent reality, the fact that ‘the expression “ a ” refers’ is true in a language does not require us to say that there is indeed some object in reality that the expression ‘ a ’ is related to – all we are doing is saying what *sentences* are true in the L language.

Remember that according to Hirsch the sentence ‘there is a table’ is true in the CSO-language and false in the N-language, and this difference in truth value has nothing to do with the world itself – the disagreement is merely verbal. Now consider that ‘ S ’ will be true in both the CSO-language and the N-language, and so ‘the expression “table” refers’ will be true in the CSO-language whilst it will be false in the N-language. However, if the original dispute over the truth of ‘there is a table’ is merely verbal, then the dispute over the truth of ‘the expression “table” refers’ will *also* be verbal; in disagreeing over the truth value of the latter sentence once again the CSO and the nihilist are not actually disagreeing over what the world itself is like. This is because a disagreement over whether an object exists and a disagreement over whether the expression that picks out that object actually refers is the *same disagreement*, and if the former is merely verbal then so is the latter. Thus it seems that we avoid the metaphysical contradiction.

That leaves us with the semantic contradiction. A more complete explanation of how the semantic contradiction can be blocked can only be fully appreciated within the context of Hirsch’s top-down semantics, and as such this issue will be returned to in section 4.2.

Let’s note first of all again the schematic sentence:

(S)If the sentence ‘there is an a ’ is true, then the expression ‘ a ’ refers

At a first glance, this doesn't appear to rid us of the semantic contradiction. Let's imagine we are speaking the N-language; 'S' is taken to be true in that language. But let us now replace 'there is an a ' in 'S' with the CSO-language sentence 'there is a table'. We now have an N-language sentence which is *about* a CSO-language sentence, and since 'there is a table' is true in the CSO-language, owing to 'S' being true in the N-language it then follows that the expression 'table' refers in the N-language, so the semantic contradiction appears to remain.

However, now consider that we have already established in section 2.2 that the same sentence can have different truth conditions in different languages owing to the different semantic rules governing sentence forms involving quantifier expressions that are operative in those languages. This is what allows us to say that 'there is a table' is true in the CSO-language but false in the N-language. Recall that for Hirsch this amounts to the claim that the quantifier expressions can be seen as functioning differently in different languages – which is the central claim of quantifier variance. He then adds that 'if we are imagining that the quantifier expressions in the M-language function differently from our A-quantifiers, then we can – and naturally will – imagine a correlative difference in the use of the word "refer". If we alter our concept of "a thing" then we alter our concept of "reference to a thing" (Hirsch, 2002: 56). What this suggests is that in each language, as there are different meanings for quantifier expressions, there is also a different meaning for the concept of reference. What does this mean in practice? In his 2009, Hirsch spells this out as accepting as true, in each language, what he calls the *weak Tarskian principle*:

(wtp) "For a sentence of the [syntactic] form ' $F(a)$ ', of *this language*, to be true, the [syntactically] singular term ' a ' must refer" (Hirsch 2009: 250) [emphasis added].

What this implies is that we can only talk disquotationally about what expressions refer to *in our home language*. Thus from the perspective of the nihilist, though he can regard 'there is a table' as a true sentence of the CSO-language, since *wtp* only requires that the syntactically singular expressions of the N-language refer, the fact that 'there is a table' is true in some other language does not then imply that the expression 'table' refers in the N-language. What this creates is a *relative* concept of reference, where each language has its own concept of reference which cannot be *used* to say what expressions of other languages refer to. It also appears to rid us of the semantic contradiction because now the nihilist can continue to regard 'there is

a table' as true in the CSO-language (owing to different semantic rules operative in that language) without being committed to a contradiction in his own language about whether the expression 'table' refers.

Let's now move on to consider another consequence arising from quantifier variance.

3.2: Translation and compositional truth-theoretic semantics

Consider that for the A-language speaker (from section 2.2), there is no way for them to say what the expression 'the thing composed of Clinton's nose and the Eiffel Tower' refers to (in their own sense of 'refer'). This is because, according to the A-language, the expression 'the thing composed of Clinton's nose and the Eiffel Tower' does not refer. This is problematic because it means that an A-language speaker cannot provide a compositional semantics below the level of the sentences³ for novel M-language sentences involving the expression 'the thing composed of Clinton's nose and the Eiffel Tower', or indeed *any* expression in the M-language which does not refer (again in the A-language sense of 'refer').

The issue here is the inability to provide a reference for the expression 'the thing composed of Clinton's nose and the Eiffel Tower' means that an A-language speaker cannot adequately specify the truth conditions for the sentence 'there exists something composed of Clinton's nose and the Eiffel Tower' in terms of the underlying reference relations of sub-sentential expressions.

This is problematic for two reasons. First of all, it is generally thought that an adequate specification of the truth-conditions of sentences in terms of the reference relations of the words in those sentences is necessary to explain how we understand novel sentences. With a finite vocabulary we have the ability to form an indefinite number of meaningful sentences. Our ability to understand such sentences is thought to depend on our knowledge of such reference relations.

³Even if the A-speakers could provide a compositional semantics in terms of the truth values of composing sentences (e.g. A & B is true iff A is true and B is true), the truth conditions of atomic sentence cannot always be compositionally analysed in terms of the semantics of the sub-sentential parts. This will become more apparent once we understand Hirsch's top-down approach to semantics.

A related and more serious worry for Hirsch is that the A-language speaker appears to be unable to translate the M-language sentences containing the expression ‘the thing composed of Clinton’s nose and the Eiffel Tower’ into sentences of his own language. There is no equivalent term in the A-language for the thing supposedly referred to by the expression ‘the thing composed of Clinton’s nose and the Eiffel Tower’. This is worrying for Hirsch because his central claim is that the A-language and the M-language are just alternative languages describing the same reality – they do not really differ in the substantive claims about what that reality contains. Thus we should reasonably expect some straight-forward way of translating between the languages to see that this is the case. However, the fact then that in the A-language we cannot say what the expression ‘the thing composed of Clinton’s nose and the Eiffel Tower’ refers to means we will be unable to translate an indefinite number of supposedly true sentences of the M-language containing that expression. This inability to translate straight-forwardly what the M-language sentence expresses into a true-sentence of the A-language would cast doubt on the idea that the M-language and the A-language ultimately are just describing the same facts in reality – it would indeed suggest that there are some truths expressible in the M-language which could not be expressed in the A-language, and thus the debate would no longer appear to be verbal. What Hirsch needs to do then is provide a way of effectively translating between the A-language and the M-language even if the A-language speaker cannot provide a compositional semantics for certain M-language sentences in terms of his own sense of ‘refer’.

How does Hirsch respond to these worries? Regarding the first worry about the inability to provide a truth-theoretic semantics for the M-language, he says if it is possible to formulate in the A-language a truth-theory for the A-language, then it is possible to formulate in the M-language a truth-theory for the M-language, each theory being formulated in terms of each language’s meaning of “reference to a thing” (Hirsch, 2002: 58). Hirsch here then is appealing to the idea that in each language it is possible to provide a truth-theoretic semantics *for that language* in terms of the reference concept *of that language*. Thus although the A-language speaker cannot say what the expression ‘the thing composed of Clinton’s nose and the Eiffel Tower’ refers to, the M-language speaker can trivially say that ‘the expression “the thing composed of Clinton’s nose and the Eiffel Tower” refers to the

thing composed of Clinton's nose and the Eiffel Tower⁴ – but of course the M-language speaker is using *his* concept of reference here, not the A-language speaker's concept. The second more problematic worry still remains however – the A-language speaker cannot translate M-language sentences containing expressions that do not refer into his own language.

In response to this worry, Hirsch says the following:

Let us note that the M-language is “translatable” into my A-language in at least the following sense: For any M-sentence I can find an A-sentence with the same truth-conditions, where two sentences have the same truth-conditions if, relative to any context of utterance, they hold true with respect to the same possible situations. This kind of intertranslatability between the A-language and the M-language holds even if it's impossible to formulate in the A-language a finitary theory of truth for the M-language (Hirsch, 2002: 57).

What Hirsch is saying here is that even though the A-language speaker cannot provide a truth theoretic semantics for the M-language he can produce in his own language a sentence that is *truth-conditionally equivalent* to the M-language sentence. Thus although the A-language speaker cannot say what the M-language expression ‘the thing composed of Clinton's nose and the Eiffel Tower’ refers to, he can say that that the M-language sentence ‘there exists something composed of Clinton's nose and the Eiffel Tower’ is true *just whenever* there exists Clinton's nose and the Eiffel Tower.

Another example to help see what Hirsch is getting at is the following: Consider that in any situation where a common-sense ontologist will say ‘there is a table’, the nihilist will say ‘there are simples arranged table-wise’. Thus, whilst the nihilist cannot provide in his own language a reference for the expression ‘table’, he can make intelligible to himself the idea that the CSO-sentence ‘there is a table’ is true just whenever there are simples arranged table-wise. Thus what Hirsch argues is that all that is necessary for translation is *the ability to assign truth-conditions at the level of sentences* – he argues that ‘The primary focus is always on whole

⁴ In this way we can see that Hirsch here is presupposing a deflationary notion of reference. For if it is true in the M-language that ‘there is an *a*’, then it would be trivially disquotational in the M-language that the expression ‘*a*’ refers to *a*.

sentences and how to assign truth conditions to them in the most charitable way possible' (Hirsch 2005: 72).

This is all that is needed, according to Hirsch, in order to be able to effectively translate between the languages. For example, he argues that someone who speaks the A-language could learn the M-language ostensively by being given a few examples of how the language works, and can then reasonably be expected to go on in the same way (Hirsch 2002: 58). For example, we could explain how the M-language worked to the A-language speaker in the following way: Whenever it is true in your language to say 'the F-thing and the G-thing exist', it is true in the M-language to say 'there is something composed of the F-thing and the G-thing'. The A-language speaker is then expected to be able to go on in the same way when faced with novel M-language sentences. Thus the M-speaker might say 'there is something composed of Big Ben and St. Paul's Cathedral', and the A-speaker, understanding how the M-language works, can then see that this sentence is truth-conditionally equivalent to the A-language sentence 'Big Ben and St. Paul's Cathedral exist'.

One might think that the issue of reference might arise here again – surely *if* the M-language sentence is true, then the expression 'the thing composed of Big Ben and St. Paul's Cathedral' refers. However, Hirsch says we are simply concerned with assigning truth conditions at the level of the *whole sentence* in such a way that the controversial sentence is seen as truth-conditionally equivalent to sentence we find non-controversial, and the sentence we find non-controversial will be a true sentence of our own language for which we *can* provide a compositional semantics (using then our own concept of reference). Once we have done that, we can effectively understand what the M-language speaker is saying in terms that are acceptable to us.

In a similar vein, Hirsch also points out how it would be easy to teach someone who speaks an endurantist-language to speak like a perdurantist: In any situation where we would describe some object as having some property *P* at *t*, the perdurantist would just say that the object 'had a temporal part which was *P* at *t*' (Hirsch 2009: 247). The point then is that even without truth-theoretic semantics it is fairly easy to switch from speaking like an endurantist to a perdurantist, or to switch from speaking like a nihilist to a common-sense ontologist. Thus the inability to provide a compositional semantics for some sentences of other

languages does not prevent us from being able to supply *a true sentence of our own language* which we can then say is *truth-conditionally equivalent* to the controversial sentence. For Hirsch then, all that is required to effectively translate between one language and the other is the ability to provide, for the sentence to be translated, a true sentence of our own language which is truth-conditionally equivalent to that sentence.

4: Top-down semantics and neo-Fregeanism in Hirsch

In sections 3.1-3.2 we saw that quantifier variance seems to require two things – a deflationary theory of reference and the idea that all that is required for translation is the finding truth-conditionally equivalent sentences in different languages, even if we cannot provide a truth-theoretic semantics for another language. In his 2009, Hirsch incorporates these implications of quantifier variance into a general approach to semantics. There are two distinctive elements to this approach to semantics: Firstly its *top-down* nature, and a particular thesis that Hirsch calls *neo-Fregeanism*. An appreciation of these elements of Hirsch's thinking is essential to understanding his idea of quantifier variance and how debates in physical object ontology are merely verbal.

To understand this position, we first need to understand Hirsch's account of sentential content and his overall conception of language. In his 2009 he explains:

I'll follow Lewis in taking a "proposition" to be a set of possible worlds. And I'll follow Kaplan in taking a sentence's "character" to be a function that assigns to the sentence, relative to a context of utterance, a proposition (the proposition being the set of worlds in which the sentence holds true). The character can also be said to give the sentences "truth conditions" (relative to a context of utterance". By the "interpretation" of a language I'll mean a function that assigns to each sentence of the language a character. Note that interpretation in this sense is defined in terms of the characters of sentences, not in terms of the reference of expressions (Hirsch, 2009: 234).

According to Hirsch then, a sentence's character, assigned to a sentence by an interpretation function (in Hirsch's sense), determines the content of a sentence relative to a context of utterance. The content in each case is a proposition, understood by Hirsch to be the set of possible worlds where the sentence holds

true. Hirsch's view of a proposition then is that it is an unstructured entity – the content of the sentence 'there is a table', in a particular context, is just the set of worlds where the sentence 'there is a table' is true. On this account, two sentences which are true at the same set of possible worlds (relative to a context of utterance) have the same content – we also say that they are *intensionally equivalent*. For Hirsch then, two different syntactically structured sentences can have the same character – for the character of a sentence is simply assigned via an interpretation function at the level of sentences, regardless of syntactic structure or underlying semantic functions of sub-sentential expressions. This particular aspect of Hirsch's view will be explained in more detail momentarily. The important implication of this though is that if two sentences are character-wise equivalent then in any context of utterance they will express the same proposition. Thus according to Hirsch, the sentence 'there is a table' in the CSO-language and the sentence 'there are simples arranged table-wise' in the N-language will have the same character C. This then determines that in any context of utterance these two sentences will express the same proposition. This view appears to stem from the idea that in any situation which is correctly describable in the CSO-language as a situation where 'there is a table' is true, that situation is also correctly describable in the N-language as a situation where 'there are simples arranged table-wise' is true. What this view of semantics amounts to then is a *top-down* view of semantics where the character of a sentence determines the content of sentence, and the character of the sentence is not determined by the underlying semantic functions of sub-sentential components. In this sense then, the character of a sentence is *explanatorily prior* to the referential functions of sub-sentential expressions. As Hirsch says:

What must be given up is a picture of language in which the characters at the level of sentences are generated by some underlying referential mechanisms at the level of words. This "bottom-up" picture is misguided because the references of words depend upon the characters of sentences (Hirsch, 2009: 248).

The idea then is that if we are looking to understand what a sentence of another language means, we should not look to *individual words* in that language but at how whole sentences are used in a particular linguistic community. Only once the character of a sentence is determined can we then say what the sub-sentential expressions refer to. And, given the deflationary theory of reference, in each case if a particular sentence 'aT' is true in *L* in context C, then it will trivially follow that

“a’ refers to a ’ is also true in L in context C . In other words, the notion of reference on this view plays no explanatory role in determining the content of a sentence.

This then is Hirsch’s commitment to what he calls neo-Fregeanism. He argues, for example, that ‘the truth values of sentences are explanatorily prior to the references of terms’ (Hirsch, 2009: 249) and additionally mentions that ‘if we are neo-Fregeans we do not explain truth conditions by appealing to reference’ (Hirsch 2009: 250). The important point to draw from this seems to be that the primary way that language relates to the world for Hirsch is at *the level of whole sentences*, not at the level of individual words⁵. This particular aspect of Hirsch’s position is important to bear in mind when discussing some of the criticisms made against Hirsch – particularly Sider’s criticisms in chapter seven. In later chapters when discussing these criticisms, the relevant implications of Hirsch’s commitment to a deflationary theory of reference, top down semantics and neo-Fregeanism will be examined in more detail.

4.1: The question of how sentences get their character and the notion of a ‘linguistic decision’

According to Hirsch, each language is a set of characters distributed over a set of syntactic structures via an interpretation function (Hirsch 2009: 249). The character of a sentence determines in each context of utterance the proposition expressed by a sentence, so ignoring for now the possibility of hyperintensional distinctions in content between sentences⁶, it could be said then that the character of a sentence supplies that sentence with its meaning. The question then though is this: In each case, what is it that determines that a specific character is assigned to a specific sentence? According to Hirsch’s top down approach to semantics, we cannot look to the meanings/references of sub-sentential expressions to show how the character at the level of a sentence is built up out of or determined by the semantic functions of individual expressions. Rather, we must look to how the sentence is deployed within a particular linguistic community – what situations is it used in and in which situations is it generally assented to. This element of Hirsch’s position will become more apparent when we discuss what I call Hirsch’s *plausibility argument* in sections 6-6.3, but for now that Hirsch seems to commit to

⁵ Indeed, an important element of being a neo-Fregean is Eklund’s claim that ‘we do not epistemically access the world by relating individual signs to objects, but by relating whole sentences or thoughts to the world’ (Eklund 2006, 97-8).

⁶ But see the later chapters on Hawthorne and hyperintensionality (chapters 5 & 6).

a meta-semantics of use at the level of sentences can be seen in the following remarks:

The primary focus is always on whole sentences and how to assign truth conditions to them in the most charitable way possible. When I speak throughout this paper about interpreting a language this is always to be understood in the narrow sense of assigning truth conditions. (Hirsch, 2005: 72)

We are looking for *genuinely plausible* truth-condition assignments that make the most charitable sense of what members of a community say, and especially the most charitable sense of how what they say rationally reflects their perceptions (hence their causal connections to their environments) and their understanding (Hirsch, 2005: 73)

The idea then is that in interpreting a language we should do so in the ‘narrow’ sense of assigning truth conditions to sentences (relative to a context of utterance). In addition to this then the assignment must make sense of the verbal behaviour of the linguistic community we are interpreting. Thus, for example, if we notice that in any situation where we utter ‘there is a table’ a member of another linguistic community is disposed to utter ‘there are simples arranged table-wise’ it would be plausible to assign the latter sentence the same character as the former. In that sense we look to how *whole sentences* are *used* in a linguistic community and assign truth-conditions to those sentences in a way that best reflects the verbal behaviour of that community. The idea that characters at the level of sentences appear to be determined by how that sentence is used within a particular linguistic community is thus another important aspect of Hirsch’s position that we must bear in mind when considering criticisms made against Hirsch.

At this point I want to clarify the notion of a *linguistic decision*, as this will be of some importance in the following two chapters. When Hirsch speaks of a linguistic decision, he means the decision to speak a particular language. Remember for Hirsch, a language is a set of characters assigned to a set of syntactic structures. Thus if we are speaking the A-language, we can make the decision to switch to speaking the M-language. Additionally this then means adopting a different set of semantic rules governing sentence forms involving the quantifier. In this sense of a linguistic decision, we are making the decision to switch from one *interpreted* language to another.

But there is also another sense in which *linguistic decision* may be used. For we may say that in each case where a character is *originally* assigned to a sentence, that is the result of the decision to *use* that sentence in particular way. Thus it is our decision to use a sentence in particular way that determines its character and hence intension in a context. Quite obviously though if we are speaking ordinary English we cannot decide to change the semantic rules of ordinary English by making the decision to speak a different way; when we learn a natural language these sorts of linguistic decisions have, as it were, *already been made*. Rather, we should say that if we are *already* speaking language *L*, we do not choose which characters are assigned to which syntactic structures, but that the meta-semantics facts which determine what characters are assigned to which syntactic structures will be facts about how sentences are *used* in the *L*-speaking community.

4.2: Bringing together Hirsch's thoughts

Now we are in position to see how the notion of a deflationary theory of reference and the lack of a requirement to provide a compositional truth-theoretic semantics for true sentences of other languages properly tie in with Hirsch's commitment to a top-down, neo-Fregean approach to semantics. Let's remind ourselves that for Hirsch, a language is a set of syntactic structures and a set of characters, such that characters are assigned to sentences by an interpretation function. It should be noted that in this section I do a great deal of explanatory leg-work for Hirsch which he himself has not undertaken in any of his own writings – he indeed neglects to draw many of these different elements of his thought together into a cohesive picture.

- 1) Note first of all that when one is speaking a language, one is *already* speaking an interpreted language, such that the sentences one uses in that language already have a character.
- 2) Secondly, note that the character, insofar as it provides in every context of utterance the intension of a sentence, can be said to provide the *content* of the sentence⁷.
- 3) Thirdly, owing to the neo-Fregean thesis, the character of a sentence is in no way constrained by the semantic properties of the sub-sentential components of that sentence – i.e. the character of a sentence is not determined compositionally. Whilst Hirsch argues that an 'interpretation

⁷ But see chapters five and six concerning the notion of hyperintensionality.

function' assigns characters to sentences, I also discussed that the meta-semantic facts which seem to determine the character of a sentence in a language are facts about how that *whole sentence is used* in a particular language.

Finally then, let us note that it is possible for different syntactic structures to have the *same character*, and thus to express the same *content* (see point two). This is possible because characters are assigned to whole sentences (syntactic structures), and this assignment is determined by how sentences are used in a particular linguistic community (see point three). Thus we can have a situation where in the CSO-language, the syntactic structure 'there is a table' has character C, owing to how that sentence (syntactic structure) is used in the CSO-language. Additionally, in the N-language, the syntactic structure 'there are simples arranged table-wise' has the *same character C*, owing to how *that* sentence (syntactic structure) is used in the N-language. The sentences thus have the same *content*, insofar as in every context of utterance they express the same proposition (which is to say they are true at the same set of possible worlds). This technical formulation makes sense of the idea then that when the nihilist says 'there are simples arranged table-wise' and the CSO says 'there is a table' they are describing the same fact.

Let us note how this relates to the issue of translation. If we are an N-speaker, we want to translate what the CSO-speaker says when he utters 'there is a table' into a true sentence of our own language (or rather, this should be possible if the claim that the dispute between the CSO and the nihilist is merely verbal is true). In order to do this, we simply need to find a character-wise equivalent sentence of our own language – i.e. a sentence of our own language which is true just whenever the CSO sentence is true. Given that character determines content, we can then regard the N-sentence 'there are simples arranged table-wise' as an effective translation of the CSO-sentence 'there is a table'. Furthermore, let us consider that, as an N-speaker, we do not need to provide a compositional semantics for the CSO-sentence because the semantic properties of the sub-sentential components *do not determine the character of the sentence*. Thus even though the N-speaker cannot say, in their own sense of 'refer', what the CSO expression 'table' refers to, this doesn't prevent them from determining that the CSO-sentence 'there is a table' is character-wise equivalent to the N-sentence 'there are simples arranged table-wise'.

We might question however, in each situation, *how* the N-speaker determines which of his sentences is character-wise equivalent to a CSO-sentence of the form ‘there is an x ’ (where x is a composite object). The attraction a compositional semantics is that we can provide an account of this for novel sentences. However, as Hirsch points out we don’t need to do that – we can look at the rules governing the use of sentences of the form ‘there is an x ’ in the CSO-language, and see that whenever it is true in the N-language that ‘there are simples arranged x -wise’, a sentence of the form ‘there is an x ’ is true in the CSO-language (see section 2.2).

Let us now consider the issue of reference. I argued that Hirsch is committed to a deflationary theory of reference. Among other things, a deflationary theory of reference argues that ‘reference’ is merely a formal device and doesn’t imply anything about there being a metaphysical relation between a word and extra-linguistic reality – it is a way from switching from *using* expressions to mentioning them within a language. The idea then is that if it is true in L that ‘there is an a ’, then it trivially follows that ‘the expression “ a ” refers’ is also true in L . Now we are in a position to see that the fact that ‘there is an a ’ is true in L is due entirely to the character of that sentence in L and has nothing to do with the semantic property of the expression ‘ a ’ (owing to the neo-Fregean thesis). Thus in knowing that ‘the expression “ a ” refers’ is true, we do not get a deeper level of explanation or understanding than we initially had from knowing that ‘there is an a ’ is true.

Now consider that ‘there is an a ’ can have a character in L^* such that it is false. It thus follows that ‘the expression “ a ” refers’ is also false in L^* . Thus a permutation in the character of the former sentence entails a permutation in the character of the latter sentence. Now consider that when an L^* speaker is talking about his own language he is *using* a certain set of syntactic structures that are already assigned characters in the L^* language. However, it is not contradictory from that standpoint to imagine that the syntactic structure S , which happens to have character C in my language, could have had a different character C^* . If we imagine that syntactic structure S does indeed have a different character than the one it does (from our own perspective) we can imagine a correlative shift in the characters of related sentences. Thus we can imagine that if ‘there is an a ’ has a different character, we can also imagine a correlative shift in the character of the sentence ‘the expression “ a ” refers’. We are now in a position to see more fully how Hirsch can avoid the semantic contradiction I drew attention to in section 3.1.

If I am an L^* speaker, then the sentence ‘there is an a ’ has a character C in my language, such that it is false in my language. As such, the sentence ‘the expression “ a ” refers’ is also false in my language. When I say in my L^* - language that the L -language sentence ‘there is an a ’ is true, I am saying that it has a different character, C^* , in that language, such that it is true. This doesn’t then change the fact that the sentence ‘the expression “ a ” refers’ is still false in *my language* owing to the character it has in my language. Thus we see that the characters of sentences in another *language do not have an effect on the character of sentences in my language*. Additionally then, when I say that ‘the expression “ a ” refers’ is true in the L language, I am not contradicting myself, because that sentence has a different character in the L language than it does in my language. Thus, from the L^* perspective, I can regard the sentence ‘there is an a ’ as being true in the L -language, whilst maintaining that ‘the expression “ a ” refers’ is false in my language. It is important not to fall into use mention confusion here. When I talk about what expressions refer in my own language, I am *using* sentences such that they have the characters they do indeed have in my own language. When I talk about other languages, I do not talk about what sub-sentential expressions refer to in that language – I only talk about what *sentences are true in that language* owing to the characters that sentences have in that language (which again has nothing to do with the sub-sentential components of the sentence).

A deflationary theory of reference is crucial in order for this to work, as I mentioned in section 3.1. It can’t be the case that the fact that ‘the expression “ a ” refers’ is true in L implies that there is some metaphysical link between the expression “ a ” and some singular object in language independent reality, because the falsity of the same sentence in another language would imply that this metaphysical link didn’t exist – indeed because it would imply that the object in language independent reality didn’t exist, which would indeed be contradictory because it implies *that in reality itself* there both is and isn’t an object which is the referent of the expression “ a ” – i.e. a deflationary theory of reference guards us against the metaphysical contradiction I drew attention to in section 3.1.

4.3: More points of contention

At this point in the discussion we can again raise additional points of contention that will form the basis for developing criticisms of Hirsch’s position in later chapters. First of all, in section 4 it was explained how Hirsch’s position involves a

commitment to a neo-Fregeanism, top-down semantics. On the approach presented, the content of a sentence is a proposition taken to be a set of possible worlds, and thus two sentences that are true at the same set of possible worlds have the same content. In chapter five, we see that Hawthorne calls this an ‘intension-centric’ approach to semantics, and he criticises Hirsch’s use of it, for it suggests that different ontological theories countenance the same set of possible worlds. Hawthorne wants to show, among other things, that ontological theories entail a smaller or larger range of metaphysical possibilities and as such not all ontologies countenance the same set of possible worlds. In addition to this, Hawthorne raises the issue of finer-grained distinctions in content. Leading on from this then, in chapter six I consider the notion of hyperintensional distinctions between intensionally equivalent sentences and show how this leads to what I have called the H-problem for Hirsch.

Additionally, we have also seen in section 4.1 evidence to suggest that Hirsch thinks that how a sentence is used within a linguistic community determines its character. This issue once again is considered by Sider who argues that there are external constraints that determine what certain expressions can mean. In chapter seven however I argue that Sider ultimately hasn’t taken into account Hirsch’s commitment to top-down semantics and neo-Fregeanism, and so many of his points will be seen to miss the mark. By contrast, in chapter eight I develop an argument against Hirsch which does take these aspects of his position into account to show that there can still be room for ontological disagreement. For now let us continue with our discussion of quantifier variance.

5: Quantifier variance and the shallow approach to ontology

By now then we have seen that quantifier variance is essentially the idea that there are different semantic rules associated with sentence forms containing quantificational expressions. The implication of this is that a sentence of the form ‘ $\exists xFx$ ’ can be true or false depending on which rules are operative in the language in which the sentence is uttered. Accepting quantifier variance, we have seen, involves a commitment to a deflationary theory of reference. Additionally, Hirsch has couched the notion of quantifier variance within the context of a neo-Fregean, top-down semantics. His overall aim in developing this position has been to show that the sort of debates covered in physical object ontology are merely verbal. He has done this by showing that a disputed sentence can have different truth-values

in different languages owing only to the rules of that language, thus showing that the disputants are not actually disagreeing about how the world itself actually is.

What this leads to, according to Hirsch is a *shallow* approach to ontology where the answers to ontological questions are trivially arrived at simply by investigating the concept of existence (i.e. the semantic rules governing sentence forms containing quantifier expressions) operative in one's home language. Hirsch argues that ordinary English is in line with the common-sense conception of ontology – thus he argues that 'once the possibility of quantifier variance is accepted...there is nothing to inhibit us from simply expressing the trivial common sense truth in terms of the quantifier we actually have in our language' (Hirsch 2002: 62). A related point he then makes is that revisionary ontologists, in making their revisionary ontological claims, are in effect making a verbal mistake, because they intend their assertions to be true in ordinary English but they are only true in some alternative language in which the relevant concept of existence is operative. As such, Hirsch says

I argue that the revisionists are in effect merely choosing to use a language different from our ordinary language, and insofar as they are not aware of this (and take themselves to be using ordinary language), they are making a certain kind of verbal mistake (Hirsch 2009: 232).

The idea then is that the sentence 'there is something composed of Clinton's nose and the Eiffel Tower' is *not true* in ordinary English. In order to know this, we investigate not the world but the semantic rules operative in English which govern sentence forms containing quantifiers. Additionally we note then that there is some language, the M-language, where the sentence is true owing to the rules of that language. Insofar as the person asserting this sentence as true claims to be speaking ordinary English, they are mistaken to the extent that they are *not* speaking ordinary English but are in fact speaking the M-language, and in that respect they have made a verbal mistake⁸. Crucially however, they are *not* mistaken about how the world itself actually is, because in the M-language the sentence 'there is something composed of Clinton's nose and the Eiffel Tower' is a true description of (some aspect of) reality. Thus it is for this reason that Hirsch argues that debates in physical object ontology are trivial and not substantive.

⁸ We can relate this back to the notion of a linguistic decision discussed in section 4.1: Hirsch is essentially arguing that the revisionary ontologist has made a linguistic decision to speak a particular way, but insofar as they are unaware of this they are making a verbal mistake in taking themselves to be speaking the truth in ordinary English.

6: *Why accept quantifier variance?*

So far we have seen what quantifier variance is and how it might be used to motivate a shallow approach to ontological questions and a trivial method for finding answers. However, it will not have escaped the reader's attention that we have yet to be given a good reason for *why* we should accept quantifier variance in the first place. In his 2002 paper, Hirsch simply introduced the notion (which he credits to Putnam) and from there shows how it may be used in his deflationary project. However, we may be entrenched metaphysical realists such that we believe there is a metaphysically privileged sense of the quantifier, and any serious discussions of ontology must be undertaken using that privileged quantifier (Hirsch 2002: 62). As such, people who are inclined to take ontology seriously are likely to simply reject quantifier variance if they have no countervailing reason to accept it. In fact, we might characterise the difference between shallow and deep ontology as turning on whether quantifier variance is accepted. Those who accept it will favour a shallow approach to ontology whilst those who reject it will favour a deep theoretical approach to ontology. Thus Hirsch has yet to provide a reason for why any committed metaphysical realist will give up on the idea of a metaphysically privileged quantifier and embrace a shallow approach to ontology.

In this section then I am going to discuss the reasons Hirsch does present for thinking his approach to ontology might be the right one. In Hirsch's writing I can detect three arguments overall, although only one is given serious, extended consideration – an argument I have termed the *plausibility argument*. Before considering the plausibility argument, I will discuss the two lesser reasons for why one might give up on the notion of a metaphysically privileged quantifier.

The first argument is that the fact that ontologists are not really disagreeing about facts is plainly obvious:

Look at your hand while you are clenching it, and ask yourself whether some object called a fist has come into existence. As shallow ontologists the first thought that must come to mind when we ask this question is this: *There can't be anything deep or theoretical here*. The facts are, so to speak, right in front of our eyes (Hirsch 2002: 67)

What Hirsch is suggesting here is that the kind of things ontologists are arguing about are not genuine issues, and that this is obvious – we just have to *look* and

see, for example, that there is no difference between claiming ‘there is a table’ and claiming ‘there are simples arranged table-wise’. Clearly this is simply a matter of temperament – those who are generally inclined to common-sense solutions will be attracted to ontological shallowness, whereas those who think such issues really do run deep into the nature of the universe will see this attitude of obviousness as simply hand-waving away difficult questions about reality.

The second argument we might call the intractability argument. This is the suggestion that ontological disputes have reached an interminable stalemate – the ‘all is said and done’ stage (Hirsch 2005: 82). This is the point where no further evidence or argument can be brought to bear on the dispute, and each side is confident in their position. The problem with a debate that reaches this stage is that it might suggest that the debate itself is ill-formulated. Thus, accepting that ontologists have all this time just been involved in a verbal debate provides an explanation for the intractability of the debate itself. If we want to move on past this intractability, we should accept this, and thus we should accept quantifier variance. Although Hirsch doesn’t explicitly argue this point himself, it places him in the company of the logical positivists and Wittgenstein who tried to show how seemingly impassable philosophical debates could be dissolved (NB: not solved) by attending properly to language. The intractability of the debate suggests that ontologists have been led down a philosophical dead end by construing their disagreement as a substantive one about the world. In accepting quantifier variance then, we have a way of making philosophical progress by dispelling mistaken assumptions about language. Again this is not persuasive for everyone – seeming intractability for some indicates not an ill-formulated debate but rather shows how difficult the right answer is to arrive at.

6.1: The plausibility argument part I

The most sustained line of reasoning I can draw out from Hirsch regarding why we should accept quantifier variance is what I call the plausibility argument. The plausibility argument is a reconstruction of Hirsch’s reasoning, found mainly in his 2005 article, for why we should accept that in some circumstances, certain principles of linguistic interpretation mean it is more plausible that someone is speaking the truth in their own language rather than speaking falsely in our language. If we accept this line of reasoning for revisionary ontological claims as well, then we should accept quantifier variance since it is required in order to

provide the non-standard interpretations of revisionary ontological claims where they come out true. Let us now look at the argument in more detail.

Imagine in some context, *C*, *A* utters 'S', which on a standard interpretation is obviously false. By a standard interpretation here I mean that if *A* is speaking the interpreter's home language *L*, then in *L*, 'S' has character *c*, such that at *C*, 'S' is false. Thus the idea of a standard interpretation is relative to the language of the *interpreter's* linguistic community. Thus we have *A* uttering 'S' at *C*, which is obviously false in *L*. We, as the interpreter, are now faced with an issue of linguistic interpretation – we need to explain the verbal behaviour of *A*, and there seem to be three options presented to us:

- 1) *A* has made a genuine factual mistake about the world itself – that is why he has uttered 'S' at *C* even though it is plainly false at *C* in *L*.
- 2) *A* has made no factual mistake about the world itself but as an *L* speaker has failed to correctly understand the language, such that at *C* he has incorrectly uttered the wrong sentence. Thus although he makes no factual mistake he has made a verbal one.
- 3) *A* has made no factual mistake about the world itself and is in fact *not speaking L*, but is instead speaking *L**. In *L**, 'S' has a character *c**, such that at *C*, 'S' is *true* in *L**. Thus the syntactic structure 'S' has a different character in *L* and *L**.

In different situations, 1, 2 and 3 could each be a plausible explanation of *A's* linguistic behaviour. What Hirsch argues is that there are three principles of linguistic interpretation which he derives from the more general *principle of charity*⁹, which indicate when it is that 3 would be the most plausible explanation of *A's* linguistic behaviour. The three principles that Hirsch talks about are the following:

Charity to perception: According to Hirsch 'any language contains sentences used to make perceptual reports, and [...] these reports are generally accurate (to a fair degree of approximation), especially when they are widely accepted in the community' (Hirsch, 2005: 71). Community-wide assent to perceptual report 'S' in context *C* suggests that 'S' is the right kind of sentence to convey the relevant

⁹ The principle of charity is discussed extensively by Davidson, but we might sum up his views on the principle of charity as 'deciding in favour reinterpreting words in order to preserve a reasonable theory of belief' (Davidson, 1984: 196).

information at C in the language of that community. Thus it is plausible that 'S' is true in *L* at C if 'S' is a perceptual report which is widely assented to by speakers of *L*. Consider then that A utters 'S' at C as a perceptual report which is obviously false on a standard interpretation, but also consider that A's assertion is widely assented to by his own linguistic community. As interpreter we must consider whether it is more plausible that A is perceptually mistaken, and thus falsely utters 'S', or that it is more plausible that A is not perceptually mistaken, but that 'S', in his own community, has a different character such that it is a true perceptual report at C.

Charity to understanding: Hirsch claims that 'there must be the strong presumption that typical speakers of a language have a sufficiently adequate grasp of their linguistic and conceptual resources so that they don't generally make a priori (conceptually) false assertions, especially when these assertions seem to be relatively simple, not ostensibly involving any complicated calculation of computations' (Hirsch, 2005: 72). This is the idea that competent language users will generally not make mistakes involving simple definitions and concepts in that language. For example, if A understands the word 'bachelor', and knows that John is an unmarried male, they will not then assert 'but he is not a bachelor'. In such a case it is plausible that A is either not fully competent in the language to know what 'bachelor' means or they are using the word 'bachelor' in a non-standard way. Again, if A appears to be from a different linguistic community and members of that community generally assent to his assertion 'but John is not a bachelor', then it is more plausible that in that community 'bachelor' means something different than what it does in ordinary English¹⁰.

Charity to retraction: Hirsch states that 'a plausibly charitable interpretation must take account of the strong presumption that reasonable people are expected to improve the accuracy of their judgements in the face of additional evidence' (Hirsch, 2005: 74). Additionally Hirsch argues that when we are trying to interpret what someone has said, we need to bear in mind 'not just what they actually have said but what they will or would say in the face of additional evidence' (Hirsch, 2005: 73). This captures the basic idea that reasonable people will tend to revise or retract their assertions if faced with additional evidence that suggests that their

¹⁰ We can relate this to the discussion of linguistic decisions in section 4.1. It would seem like a plausible constraint on assigning characters to syntactic structures that we do not thereby end up with contradictions in the language. Thus a linguistic community would not speak a language (i.e. use sentences in a certain way) that contained basic contradictions.

initial assertion was mistaken, arising from the basic desire to increase the number of true beliefs one has and to correct any false beliefs. Thus if A asserts 'S' at C, and 'S' is obviously false on a standard interpretation, we need to consider whether A would revise his assertion if we present him with further evidence or further explanation of the meaning of terms. A refusal to retract the claim after such additional evidence/explanation suggests then that A is using 'S' non-standardly. Once again if A is from a different linguistic community and they assent to A's assertion, even in the face of additional evidence and explanation then it is more plausible to interpret 'S' as having a different meaning in that linguistic community than it does in our own.

Overall, Hirsch argues that when we are interpreting what someone says, 'We are looking for *genuinely plausible* truth-condition assignments that make the most charitable sense of what members of a community say, and especially the most charitable sense of how what they say rationally reflects their perceptions (hence their causal connections to their environments) and their understanding' (Hirsch, 2005: 73). Thus an interpretation that does not make charitable sense in this way looks like an implausible interpretation according to Hirsch.

Let's now look at an example where these principles of linguistic interpretation come into play such that 3 becomes the most plausible explanation of someone's linguistic behaviour, as opposed to 1 or 2. Hirsch's asks us to imagine a community of speakers, the A-community, who speak a language very much like ordinary English, except that they will assert 'all glasses are cups', and will, when presented with a standard drinking glass, assert 'this (here) is a cup' (Hirsch, 2005: 69-70). Let's imagine some situation S then when a typical A-speaker comes along and is presented with a glass and says 'this (here) is a cup'. We are now faced with an issue of interpretation, and we need to consider whether the A-speaker has made some genuine factual mistake about what is in front of his face, whether he doesn't understand the English words 'glass' and 'cup' properly, or whether he is speaking A-English, such that in A-English 'this (here) is a cup' is true at S. Hirsch thinks that the latter option is the most plausible explanation of the A-speaker's linguistic behaviour. For example, he asks 'is it possible that, contrary to all evidence, the members of the A-community really mean by 'cup' what we mean, but they have for some reason the intractable inclination to falsely judge that glasses are cups?' (Hirsch, 2005: 70-1). He continues, suggesting that 'if we tried to interpret the word 'cup' as meaning in the A-language what it means in our

language, we would have to depict the A-speakers as inexplicably making false and unreasonable judgments about cups' (Hirsch, 2005: 71). The idea then is that if we insist that the A-speaker means by 'cup' what we mean, then this does not best reflect their causal relation to the environment or their understanding¹¹ (Hirsch 2005: 73).

For example, we would have to regard the A-speaker as being constantly perceptually mistaken whenever presented with a glass if he asserted always 'this (here) is a cup'. The fact that this perceptual report is assented to in his own linguistic community suggests then that he is not perceptually mistaken and that 'this (here) is a cup' is a true perceptual report in his language when presented with a glass. Furthermore, if we assumed that in the A-language that 'cup' has the same meaning as it does in our language, then it would look like A-speakers would be making a priori false judgements such as 'a glass is a cup' (Hirsch, 2005: 72). Thus charity to understanding suggests it is more plausible to interpret 'cup' differently in the A-language. Finally, that the A-speakers refuse to retract their assertions suggests that they must mean something different by 'cup' than what we do. All this points to the plausible interpretation that A-English is *not* ordinary English, and that when the A-speaker asserts 'this (here) is a cup' at S, he is not making a genuine factual mistake about the world, nor is he being an incompetent English speaker – rather he is speaking the A-language, and in the A-language 'this (here) is a cup' is true at S.

An important aside to note at this point is the relevance of the existence of the imagined A-community. If there was no imagined A-community which assented to the A-speaker's assertions, or if indeed the A-speaker claimed to speaking ordinary English, then it is more likely that the A-speaker is in some sense attempting to speak ordinary English (not A-English) and is thus making a verbal mistake. This is similar to Hirsch's claim that revisionary ontologists, in thinking that their claims are true in ordinary English, have made a kind of verbal mistake (but not a genuine factual one about the world itself) – see section 5.

As such then we have seen that there are times when the principles of linguistic interpretation outlined by Hirsch mean that it is more plausible to interpret someone as speaking the truth in a different language rather than to assume they

¹¹ By understanding here, I think Hirsch means a speaker's general grasp of relatively simplistic concepts of language and basic reasoning skills involving those concepts.

are speaking falsely in ordinary English. The next step is to find a plausible truth-condition assignment to the disputed sentence which best reflects the speaker's casual relation to their environment and their understanding. With the A-community, this is a relatively straightforward affair for we can simply regard the word 'cup' as having a different meaning in the A-community, and this then explains their linguistic behaviour in a way that rationally reflects their causal relation to their environment and their understanding. Let's now look at this line of argument applied to ontological claims.

6.2: The plausibility argument part II

Hirsch argues that if we are to assign truth-conditions to revisionary ontological claims that make them come out true in an alternative language in a way that rationally reflects the causal relation to the environment of the speakers of that language and their understanding then we will need the resources of quantifier variance. Although Hirsch doesn't specifically state that we have to accept quantifier variance here, he assumes that if one is in agreement that the most plausible explanation of a revisionary ontologist's assertion is that they speak the truth in their own language (and we should accept this if we accept the applicability of the principles of linguistic interpretation to ontological claims) then because we can make use of the resources of quantifier variance then we should. Let's then look at an example of how this works.

Imagine in some situation S^* where there is a table in a room, the N-speaker says 'there is no table (here)', and in general says 'there are no tables (anywhere)'. We are now faced with an issue of interpretation. Again let's assume the N-speaker is part of a linguistic community that generally speaks the same way he does. Let's say they assent to the N-speakers assertion at S^* . Charity to perception tells us that it is not plausible that this entire community has an incurable tendency to not see tables – this does not reflect their causal relation to their environment as they tend to avoid bumping into tables etc. Additionally, the N-speaker appears to understand the definition of 'table', but still refuses to retract his assertion that 'there is no table (here)' at S^* . Clearly then, the linguistic principles discussed by Hirsch mean that the most plausible explanation of the N-speakers linguistic behaviour is that he speaks the truth in his own language (the N-language) and that at S^* , 'there is no table (here)' is true in the N-language.

That decided, we must now decide how to assign truth-conditions to the N-language sentence ‘there is a no table (here)’ in a way that rationally reflects their causal relation to their environment and their understanding. One initial suggestion might be to reinterpret the predicate ‘table’ in the N-language in the way we reinterpreted the word ‘cup’ in the A-language, but we quickly see that this will not work. This is because the N-speaker will accept the following claim “For all x , if x is a piece of furniture with a flat top and one or more legs, then x is a table”. That is, the N-speaker will accept the standard definition of ‘table’. Instead, what N doesn’t accept is that there are any objects in the universe that the predicate ‘...is a table’ applies to. For that reason we can’t say that the N-language expression ‘table’ means ‘...’ and fill in the blank with an ordinary English predicate.

The other suggestion requires making use of quantifier variance, and thus if we want to assign truth-conditions to the assertion ‘there is a no table (here)’ at S^* in a way that rationally reflects the N-speakers causal connection to his environment and his understanding, then we should accept quantifier variance. In this situation, if we regard the N-language quantifier as being semantically restricted to range over only simples, as Hirsch suggests we can do¹². If we do this, then we can interpret the N-sentence ‘there are no tables (anywhere)’ as being intensionally equivalent to ‘there is no simple which is a table (anywhere)’, which is a true sentence in our own language. And thus at S^* we can interpret ‘there is no table (here)’ to mean ‘there is no simple which is a table (here)’, with the latter being true in our own language. Thus by understanding the N-language quantifier as restricted to range over only simples, we can understand why they say ‘there is no table (here)’ yet still manage to not walk into the table. This is because they do accept ‘there are simples arranged table-wise (here)’ as being true at S^* , and thus we can regard this as being intensionally equivalent to our own ‘there is a table (here)’, which is also true at S^* . In this way we can make rational sense of the N-speakers linguistic behaviour, but in order to do so we needed to be able to interpret their quantifier as being semantically restricted to range over simples.

¹² He argues in general that the side that appears to accept more objects in their ontology can make charitable sense of the assertions of a side that appears to reject the existence of those objects by understanding that their quantifier is semantically restricted to range over only certain types of entity (see Hirsch, 2005: 76).

6.3: Some comments on the plausibility argument

The most sustained line of reasoning for accepting quantifier variance to be found in Hirsch then seems to be an appeal to a more general principle of linguistic interpretation – the principle of charity; the idea that it is, on the whole, more plausible to reinterpret the meaning of an assertion in order to maintain a reasonable theory of belief. It seems then that for Hirsch, it is simply implausible to suggest that when the nihilist asserts ‘there are no tables’ that in fact his beliefs about the world and what it contains are so radically mistaken and different from our own. This simply isn’t plausible and it is more plausible to assume that his beliefs about the world are largely in line with what we believe, and to accept that he has a different language for describing that world and his beliefs about it. Thus we should avoid interpreting the N-speaker in a way that would attribute to him plainly and obviously false beliefs. The idea of Hirsch’s argument then seems to be that because we can make so much sense of the linguistic behaviour of the N-speaker, and the M-speaker, and so on, it is not plausible to think their general picture of the world is so radically different from our own. All we have to do is reinterpret what they say to maintain this theory of belief which is much in line with our own.

Thus a further motivation seems to be the idea that an *entire linguistic community* cannot be so systematically and fundamentally mistaken about the world in front of their faces, especially when making perceptual reports, and we should therefore interpret their perceptual reports in a way that reflects a rational theory of belief about the world itself. However, we might object to this line of thinking in the following way. When the N-speaker asserts ‘there is no table (here)’, Hirsch is construing this as a straight up perceptual report about what the N-speaker is perceiving in front of his face. However, what Hirsch overlooks is that ontological claims are thought to be *highly theoretical*. Now we might argue that highly theoretical beliefs can be held to be true or false in different communities and because of the nature of the highly theoretical beliefs, these differences in beliefs are not manifested in ordinary everyday behaviour and as such these beliefs, on the whole, have no impact on how an individual causally interacts with the world at an everyday level.

Thus when the N-speaker asserts ‘there is no table here, but there are simples arranged table-wise’, he is not really making a straight up perceptual report but is

instead giving a theoretical interpretation of the world that appears to him in his perception, and trying to say something about its inner nature. To the extent that it contains perceptual elements, we could then say that even if we cannot for sure determine whether the assertion is true or false, it is at least *empirically adequate* – believing that ‘there is no table here, but there are simples arranged table-wise’ will be an empirically adequate belief to the extent that it will not have an adverse effect on your ability to make empirical predictions or successfully navigate the environment (e.g. not bumping into the table) if it turns out to be a false belief. Thus we might argue that if A believes theory T and theory T* to both be empirically adequate, but believes in addition to this that T also happens to be true, then it is not implausible for him to interpret the sentences of theory T* as being false. Additionally, it is not then implausible for him to think that someone who accepts T* holds many false beliefs and many false sentences as long as those sentences are empirically adequate. Thus if we look at things in this way, it is not at all clear why we should be compelled by the plausibility argument to accept that ontologists must speak the *truth* in their own language – all we have to do is accept that ontological theories are all empirically adequate.

7: Characterising Hirsch’s position

As I mentioned earlier, Hirsch takes his position to be inspired partly by Putnam’s internal realism, and also takes his position to have more in common with Carnap and his notion of ‘external’ and ‘internal’ questions – describing his position as ‘roughly Carnapian’ (Hirsch 2009, 31). However he considers his position to differ from Carnap’s in several important respects.

Let’s first characterise what is supposed to be the central claim that Hirsch makes regarding debates in physical object ontology. These debates meet Hirsch’s criterion of being ‘merely verbal’, in that they meet the following condition:

‘Each side can plausibly interpret the other side as speaking a language in which the latter’ asserted sentences are true’ (Hirsch, 2009: 231).

This claim, I think, is important in clarifying Hirsch’s position since it is an explicit statement of a criterion whereby a debate is merely verbal. Much of the preceding chapter has been concerned with how Hirsch establishes that this criterion applies to debates in physical object ontology. Whether this is a good criterion for judging a

debate to be merely verbal will be explored in more detail in the following chapters.

In addition to laying out this criterion, Hirsch commits himself to several important theses which serve to distinguish his position from a thoroughgoing Carnapianism. In making these commitments then, I take Hirsch to be expounding a modified or 'light' variety of Carnapianism, which attempts to avoid some of the more problematic elements of that position. These commitments are the following:

- 1) Hirsch takes himself to be a 'robust realist'. He claims that "whereas Carnap's formulation sometimes seems to suggest an anti-realist or verificationist perspective, my position is robustly realist. I take it for granted that the world and the things in it exist for the most part in complete independence of our knowledge of language" (Hirsch, 2009: 231). Thus we can take Hirsch as understanding his modified Carnapianism as implying neither a variety of anti-realism or verificationism – both of which can be ascribed to a more thoroughgoing Carnapianism.
- 2) Hirsch rejects outright a verificationist explanation of how issues in ontology reduce to linguistic choices – "if the explanation is verificationist, appealing to the idea that ontological issues are hard or impossible to resolve, I reject that explanation" (Hirsch, 2009: 231).
- 3) Hirsch envisions his semantic deflationism as applying only to "questions about the existence and identity of highly visible physical objects", whereas Carnap takes his external/internal distinction to apply to "all issues in ontology, including those involving abstract things such as sets and properties" (Hirsch, 2005: 67). Thus Hirsch does not envision his position as a global deflation of all ontological debate as reducing to merely linguistic decisions.
- 4) Finally, Hirsch argues that a central goal of his project is to vindicate common sense, whereas this was not a concern of Carnap (Hirsch, 2009: 232). It is not entirely clear in what respect Hirsch is intending to defend common sense here. If we take him to mean he defends common pre-philosophical assumptions, such as the objective reality of the world and the things in it, then Hirsch's defence of common sense can be seen to go hand in hand with a rejection of any anti-realist readings of his position.

Thus we can characterise Hirsch's modified Carnapianism in the following way:

- 1) It is robustly realist, as opposed to being anti-realist.

- 2) It rejects verificationist explanations of ontological questions reducing to linguistic choice.
- 3) It is restricted to a certain subset of ontological debate, namely physical object ontology.
- 4) It seeks to vindicate common sense.

Against this we can understand thoroughgoing Carnapianism as follows:

- 1*) It implies, or can be seen to imply, anti-realism.
- 2*) Verificationist explanations are not rejected.
- 3*) It is unrestricted in its application.
- 4*) It is not concerned to vindicate common sense.

The central claim of this thesis will be that Hirsch's modified Carnapianism is ultimately not a tenable position, as I will demonstrate in the following chapters. What this means, for meta-ontological debate, is that whilst a thoroughgoing Carnapianism might yet still pose a challenge to the claim that debates in physical object ontology are substantive, there isn't a space available between the extremes of accepting this thoroughgoing Carnapianism or accepting substantive ontological debate. It appears that in characterising his position in the above way, Hirsch has attempted to stake out a middle ground – showing that the debates in physical object ontology are merely verbal without thereby having to adopt thoroughgoing Carnapianism. To the extent that some might be sceptical of the debates of physical object ontology, whilst also being sceptical of the anti-realist/verificationist implications of thoroughgoing Carnapianism, Hirsch's middle ground presents an attractive position. However, this middle ground is not a tenable position to hold.

7.1: Summary of chapter two and general outline of the argument of the thesis

In this chapter I have introduced the meta-ontological theory of Eli Hirsch, who argues that the debates in physical object ontology are merely verbal. The people involved in these debates are not actually in any disagreement about how the world is in itself (as they take themselves to be), but are in fact unknowingly describing the same facts using different languages. This is made possible by quantifier variance, which is the thesis that there can be different but equally correct meanings for quantifier expressions. This was analysed as there being different

semantic rules operative in each language which govern sentence forms involving quantifier expressions. For this to work, we also saw that Hirsch has to endorse a deflationary theory of reference and also give up on the idea of being able to provide a compositional truth-theoretic semantics for an alternative language from the perspective of our home language. Thus we then saw that Hirsch's position also involves a commitment to neo-Fregeanism and a top-down approach to semantics. In addition to this I then discussed reasons for why we should accept quantifier variance and showed that the strongest line of reasoning to be found in Hirsch's writings can be reconstructed as what I called the plausibility argument, which ultimately amounts to the claim that it is more reasonable to suppose that a revisionary ontologist speaks a different language in which his assertions come out true rather than supposing that holds many plainly and obviously false beliefs, though in section 6.3 I offered some thoughts on why it is not always implausible to regard someone as holding many false beliefs if those beliefs are highly theoretical and empirically adequate. Finally, in section 7 I outline the main claim of this thesis. Whilst Hirsch presents his position as being a way of understanding the debates in physical object ontology as being merely verbal, he does so in a way that seeks to avoid a thoroughgoing Carnapianism. In that section I claimed that this was not a tenable position to take up. The purpose of the remainder of this thesis is to explain how.

In chapter three, I pick up on issues raised by Bennett regarding the so called 'semantic rules' operative in different languages. Bennett claims that this means that certain necessarily true sentences, such as the CSO sentence 'if there are simples arranged α -wise then there is an α ', will have to be analytic if they are merely just semantic rules, as Hirsch suggests that they are. I argue, contra Bennett, that the most pressing issue this creates is that the CSO-speaker, whilst regarding the sentence as being necessarily true, *does not regard it as being analytic*. In making this point, Bennett has hit on something important. That is the idea that there can be two languages that are phonetically and intensionally equivalent, but that this can actually mask underlying differences in ontological commitment between speaking the different languages. Thus, the difference in ontological commitment can only be borne out if we attend to the truth values of certain meta-level semantic propositions about which object-level sentences are analytic or synthetic. Thus at the end of chapter three I make the case that Hirsch cannot establish that the debate between the nihilist and the common-sense

ontologist is merely verbal without assuming that the common-sense ontologist is speaking a language that the common-sense ontologist explicitly denies they are speaking. This, I argue, is to fail to take the debate between the nihilist and common-sense ontologist on its own terms. Thus it will appear, by the end of chapter three, that Hirsch's criterion of a merely verbal debate is on shaky ground since it seems to ignore the fact that ontological disagreement will manifest itself in a dispute not merely over object-level sentences, but also in a dispute over what meta-level semantic propositions are regarded as true. I call this problem for Hirsch the *analytic problem*. In order to avoid the analytic problem, I argue that Hirsch has to give up on his claim of being anti-verificationist. Thus by the end of chapter three I will show that Hirsch cannot show that the debate between the nihilist and the common-sense ontologist is merely verbal without giving up on one claim that makes his position distinct from a thoroughgoing Carnapianism.

In chapter four I develop what I think the most appropriate response to the analytic problem would be – showing that the debate over what sentences are analytic *itself* could be analysed as a further verbal debate. What I then argue is that each side in the debate can maintain a merely nominal distinction between the necessary true object-level sentences they regard as analytic and the ones they regard as synthetic in their *own* object-level language, whilst being able to charitably interpret problematic necessarily true sentences in the opposing sides object-level language as being analytic. This appears to do away with the problem uncovered in the last chapter – that ontological disagreement manifests itself in what meta-level semantic propositions as true. However, I then argue that this actually leads Hirsch to a commitment to all necessary truths being analytic. I then argue from this commitment that Hirsch cannot maintain two important claims – firstly he cannot sensibly maintain the claim that he is a 'robust realist', and secondly he cannot maintain the claim that his position only applies to debates in physical object ontology, since a number of other ontological debates end up being merely verbal if all necessary truths are analytic. Further to this, since I take it that a centrally important aspect of common-sense is a commitment to realism, then given that Hirsch can't maintain a 'robustly realist' stance, then by proxy he cannot be said to be vindicating common sense. Thus by the end of chapter four it will be seen that the only way Hirsch can avoid the issues uncovered in chapter three will be to retreat further into thoroughgoing Carnapianism, thus further demonstrated that his middle-ground position is untenable.

In chapter five I move on to discuss two central criticisms raised by Hawthorne. The first criticism is that Hirsch supposes that different ontological positions simply re-label the same set of possibilities and are thus all intensionally equivalent, and this intensional equivalence thus explains how the debates are merely verbal. Hawthorne argues that this gets things wrong as different ontological positions suppose narrower or wider ranges of possibilities and thus aren't intensionally equivalent. I defend Hirsch from this charge, showing that Hawthorne's examples are not effective at proving their point. I then discuss the second point that Hawthorne raises – namely that ontological positions make use of hyperintensional operators that supposedly track structural features of reality, and this is in tension with how hyperintensionality would be viewed on Hirsch's approach. However, Hawthorne merely drops this as nothing more than a mere hint, not really developing this into an effective criticism of Hirsch to demonstrate why his position ultimately fails to show that the debates in physical object ontology are merely verbal. This leads me on to chapter six, where I argue that in order for the debate between the nihilist and the common-sense ontologist to be merely verbal, hyperintensional distinctions in content have to be considered as merely representational, rather than as arising from tracking features of the world that are not revealed at the level of intensional equivalence. This, I will argue, demonstrates the same problem as highlighted in chapter three – that ontological disagreement manifests not only at the level of which object-level propositions are true but also again at the meta-level – specifically here regarding the claim of whether two intensionally equivalent sentences are true of the same fact, or two distinct but necessarily co-obtaining facts. Against such an objection, I argue that Hirsch must retreat further into thoroughgoing Carnapianism. Thus chapter six serves to further demonstrate that Hirsch's middle-ground position is untenable.

In chapter seven I move on to discuss the debate between Sider and Hirsch. Sider takes issue initially with Hirsch's claim that there could be different rules of use governing quantifier expressions, thus calling into question whether quantifier variance is actually possible. This leads to a complex debate which I then trace the contours of, ultimately showing that Sider hasn't taken into consideration Hirsch's commitment to a top-down, neo-Fregan approach to semantics, which I have shown in sections 4 – 4.2 of this chapter. The result of this is that Sider doesn't take into consideration that for Hirsch, the primary way in which language relates

to the world is at the level of the sentence, not at the level of sub-sentential expressions, and this is why Sider's criticisms cannot really find their mark.

In chapter eight I thus move on to develop criticism of Hirsch which does indeed take into consideration the idea that language relates to the world at the level of the sentence. I argue this still leaves room for ontological debate because the fact that any syntactic structure can be assigned any character, and thus could be true of any fact. I show that this doesn't then preclude that the fact which the sentence is indeed true *of* from having a particular kind of ontological structure. To avoid this problem, Hirsch needs to commit to facts being unstructured entities (which he does). After exploring several options for Hirsch, I argue that the notion that facts are unstructured can only be made sense of by adopting one of three views: First, he can accept that the world, itself, as a single concrete entity, cannot have any objectual structure of the kind needed to support genuine *de re* modal claims. This allows Hirsch to maintain a commitment to realism, but it also then means Hirsch is committed to monism. This position is incoherent since monism is a first-order position which is supposed to rival the positions of nihilism, universalism etc. The second position that Hirsch can adopt is to retreat further into thoroughgoing Carnapianism, so once again abandoning the middle ground. The final position, which avoids monism and a further retreat into Carnapianism, is to advocate a Tractarian inspired view whereby the world consists of the totality of primitive unstructured facts. I argue that this position, if it is to be held alongside the claim that the debates in physical object ontology are merely verbal, collapses into obscurity. Thus Hirsch, by the end of chapter eight, will be seen to be forced into incoherence, obscurity or will have to give up his middle-ground position and adopt a thoroughgoing Carnapianism.

Ultimately then it will be shown that Hirsch's modified light variety of Carnapianism cannot be tenable, and thus one is forced to choose between accepting that there is a substantive debate to be had regarding the ontology of highly visible objects, or accepting a more thoroughgoing Carnapianism.

Chapter three: Bennett and the *analytic problem*

1: *Introduction*

In chapter two I introduced Hirsch's reasons for thinking that debates in physical object ontology are merely verbal. On his view, a sentence such as 'there is a table' can be false in one language, whilst speakers of that language can charitably regard someone who asserts 'there is a table' as coming out true in their own language. In her 2009, Bennett points out that not only will a sentence like 'there is a table' be false in the N-language, the following sentence will also be necessarily false; 'if there are simples arranged table-wise then there is table'. This sentence Bennett calls a *linking principle* since it links things that are countenanced by the nihilist (simples) to things that are countenanced only by the CSO (who countenances both simples and composites). Bennett argues then that the sentence; 'if there are simples arranged table-wise then there is table' (from now on shortened to 'LP' to highlight that it is an example of a linking principle) will be necessarily true in the CSO-language. Thus the nihilist will not only have to regard the sentence 'there is a table' as being true in the CSO-language, they will also have to regard 'LP' as coming out as necessarily true in the CSO-language.

In this chapter I will first discuss a preliminary worry that Bennett raises against Hirsch, arguing that this does not constitute a serious problem for Hirsch. I will then analyse Bennett's first major contention – that for the debate between the nihilist and the CSO speaker to be verbal, 'LP' will have to be not only necessarily true, but also analytically true. I will fill in two explanatory gaps in this first contention which Bennett leaves out – first exactly what it could mean to say that a sentence like 'LP' is analytic, and secondly *why* it would have to be analytic. Surprisingly, Bennett just assumes this point without explaining why. The reason, I argue, is that if 'LP' is analytic, the N-speaker can charitably regard it as being true in the CSO-language owing to nothing more than the rules operative in the CSO-language. They can then go on to regard it as false in their own language since the CSO-rules which make 'LP' analytically true in the CSO-language are not operative in the N-language.

Bennett then makes her second main contention – that 'LP' cannot be analytic in the CSO-language. She presents two main reasons for thinking this. The first is that if we make 'LP' analytic, we ignore the *numerical distinctness* between the referents of the expressions 'simples arranged table-wise' and 'table'. Bennett then

argues that if we do take into account this numerical distinctness claim, the fact that 'LP' is analytic implies that we can define things into existence – a result she rightly supposes is absurd. Against this point, I argue that Bennett has not paid sufficient attention to Hirsch's top-down, neo-Fregean approach to language and his commitment to a deflationary theory of reference – which I explored in the previous chapter (sections 3 – 4.2). As such I argue that the first reason for why 'LP' couldn't be analytic isn't as problematic as Bennett supposes.

The second reason for thinking that 'LP' cannot be analytic is underplayed by Bennett to the extent that she doesn't give it much treatment, but here I argue that this reason is actually far more problematic for Hirsch. That reason is that the CSO speakers themselves do not regard 'LP' as being analytic – instead they think 'LP' describes some necessary feature of the world itself. In section 6 I argue how this is a serious problem for Hirsch's central contention. This is because his argument for how the debate between the nihilist and the CSO is merely verbal ignores the ways in which differences in ontological commitment manifest themselves not only in what object-level sentences one asserts, but *also in what meta-level semantic propositions one asserts*. I thus argue that if the disputants in an ontological dispute have a stake in the truth-value of some meta-level semantic proposition, Hirsch cannot assume the truth-value of that proposition without begging the question against at least one side in the ontological debate. In order to avoid this criticism, I argue that Hirsch can only retreat into thoroughgoing Carnapianism by abandoning his commitment to being anti-verificationist (see section 7 of the previous chapter). Thus the ultimate lesson to be drawn from this chapter is that the analytic problem, as developed in section 6 of this chapter can only be overcome by Hirsch if he retreats from his modified light version of Carnapianism. This supports my central contention that Hirsch's supposed middle-ground, whereby he claims to be able to demonstrate that the debates in physical object ontology are merely verbal without giving in to the anti-realist/verificationist tendencies of a thoroughgoing Carnapianism, is not a tenable position.

2: A preliminary worry: What are the sufficient conditions for a merely verbal debate?

One of the first things that Bennett points out in her 2009 is that the kind of debates that are had in physical object ontology do not appear to fulfil the sufficient

conditions that Hirsch himself sets out in order for a debate to be verbal. Whilst Bennett herself thinks this is a potentially serious issue for Hirsch (albeit one she appears to give him a pass on), I will argue that a proper regard of Hirsch's views on language will show that this initial criticism is not really problematic for him.

In his 2005, Hirsch sets out what he calls the paradigm case of a verbal dispute:

The simplest paradigm of a verbal dispute – the simplest way it can happen that each side of a dispute can find a charitable interpretation that makes the other side come out right – is where, for each disputed sentence *D*, there are two undisputed sentences *U1* and *U2*, one true and one false, such that one side holds that *D* is (a priori necessarily) equivalent to *U1* and the other side holds that *D* is equivalent to *U2*. Each side can then assign charitable truth conditions to *D* in the other side's language simply by assuming that in that language the other side's asserted equivalence holds. (Hirsch, 2005: 83)

Bennett constructs a simple example where the conditions above are met and the debate then obviously becomes verbal. Imagine we have the following:

U1: 'There is a mixture of alcohol¹³ in a V-shaped glass on the table'

U2: 'There is a mixture of gin/vodka and vermouth in a V-shaped glass on the table'

D: 'There is a martini on the table'

So *U1* in our example is undisputedly true, *U2* is undisputedly false and *D* is disputed – one side will claim it is true and another side will claim it is false. Imagine then we have some old school drinkers, who I will call the OS's, and we have some modern bartenders, who we will call the MB's. Bennett argues that an OS will regard *D* as being a priori necessarily equivalent to *U2*, and so will regard *D* as false on this particular occasion, whereas an MB will regard *D* as being a priori necessarily equivalent to *U1* and so on this occasion will regard *D* as being true (Bennett, 2009: 51).

It is important to note that both sides are in agreement on the *facts* – since they both accept *U1* as true. What they disagree over is the truth-value of *D* on this

¹³ In order for this example to work, the alcoholic mixture here has to be something *other than a mixture of gin/vodka and vermouth*.

occasion because the MB's regard *D* as equivalent to *U1* whereas the OS's regard *D* as equivalent to *U2*. This is merely a verbal disagreement because each side can charitably interpret the other side's stated equivalence as coming out true in their respective language. For example, the OS's can charitably regard *D* as being equivalent to *U1* in the MB language, whilst maintaining in their own language that *D* is equivalent to *U2*. Likewise, the MB's can charitably regard *D* as equivalent to *U2* in the OS language whilst maintaining that *D* is equivalent to *U1* in their own language. Thus, there is no real disagreement and each side is able to charitably interpret the other side as speaking the truth in their own language.

One thing to note is that such a conciliatory outcome is not always so readily available if the disputants are part of the same linguistic community due to the notion of *semantic deference*. Thus although the MB might take *D* to be equivalent to *U1*, if they are part of the same wider linguistic community as the OS, then it might be the case that semantic deference indicates that MB is wrong to regard *D* as equivalent to *U1*. For example, if wider use within the community indicates that *D* is actually equivalent to *U2*, then it would look like MB is incorrect. Furthermore, we often defer to experts regarding the meaning of certain expressions, and how they use a particular expression then determines what that expression means in the wider linguistic community. If we imagine on this particular occasion that the wider community defers to how the expression 'martini' is used by cocktail connoisseurs, then it might again be the case that the MB is wrong to regard *D* as being equivalent to *U1*. In either case though it should be stressed that the only mistake that MB is making is a purely verbal one – he is not mistaken about the non-semantic facts, and as such his disagreement with the OS would still be verbal. Despite this, the fact remains that it is possible for each side in the debate to charitably interpret the other side as speaking the truth in their own language¹⁴.

Here Bennett makes her preliminary criticism against Hirsch, arguing that the debates in physical object ontology do not fulfil the sufficient conditions Hirsch himself sets out for a verbal debate. Consider a dispute between the nihilist and the CSO. Both regard the following claim as undisputedly true:

- (1) 'There are simples arranged table-wise'

¹⁴ As Hirsch pointed out, it is a certain kind of verbal mistake to think one is speaking the language of the wider linguistic community but to actually be speaking an idiosyncratic language. This is the sense in which Hirsch regards revisionary ontologists to be making a verbal mistake (see chapter two, section 5).

In addition to this, the CSO regards the following as truth-conditionally equivalent to (1):

(2) 'There is a table'

and so will regard (2) as being true. However, the nihilist does not regard (2) as being truth-conditionally equivalent to (1) and regards it as false. However, as Bennett points out *there doesn't appear to be an undisputed falsehood that the nihilist takes (2) to be equivalent to* (Bennett 2009: 52). The reason the debate between the MB and OS looked verbal is that each side associated with the disputed sentence a different set of truth-conditions. However, in this case, it appears that the CSO associated the disputed sentence with an undisputedly true sentence, but the nihilist simply appears to deny this equivalence. As such it looks as if the dispute between the nihilist and the CSO doesn't meet the sufficient conditions set out by Hirsch.

One way around this might be to try and supply the undisputed falsehood for the nihilist, thus allowing the debate to meet the sufficient conditions. The only problem is it doesn't look like there is any suitable candidate. One suggestion might be the following. If we recall from chapter one, Merricks argued that in order for something to count as a physical object it needs to have non-redundant causal powers. We might consider the following as the undisputed falsehood then: 'There is a table (with non-redundant causal powers)'. However, the problem with this suggestion is that it is *not undisputedly false* since some will argue (Lowe for example – see chapter one) that tables *are not* causally redundant, and so would not regard 'there is a table (with non-redundant causal powers)' as false. As such, it looks as if Hirsch will have difficulty supplying the undisputed falsehood.

The inability to provide an undisputed falsehood is not seriously problematic for Hirsch if we consider that all he has to say is that the character assignments in the CSO-language are such that 'there are simples arranged table-wise' and 'there is a table' have the same character in that language, such that in every context of utterance they have the same intension, and are thus true at the same set of possible worlds. However, in the N-language, the sentences do not have the same character. So whilst the fact that we cannot say what undisputed falsehood the nihilist will regard as being equivalent to 'there is a table'¹⁵, this in itself does not

¹⁵ Another suggestion might be 'there is a simple which is a table'. The nihilist and the CSO speaker would both certainly regard this as untrue. But when the nihilist claims 'there is

seem to pose an issue for the claim that the debate is merely verbal¹⁶. Therefore let us move on to discuss Bennett's first major claim.

3: 'LP' has to be analytic in the CSO-language

The CSO argues that if there are simples arranged table-wise then there will also be a numerically distinct object, a table, which stands in relation to the simples such that the simples are part of the table. As such, whenever it is true that there are simples arranged table-wise, the CSO will claim that in the same region there is also a table. The nihilist disagrees with this and argues that although there are simples arranged table-wise, there is no additional, numerically distinct object called the table. Thus as well as disagreeing over the truth value of the sentence 'there is a table', the nihilist and the CSO will disagree over the truth value of 'LP'¹⁷. For the CSO, 'LP' is necessarily true, whereas for the nihilist, it is necessarily false.

In order for the debate between the nihilist and the CSO to be verbal, it has to be possible for the nihilist to charitably interpret the CSO's assertions as being true in the CSO-language. This includes apparent perceptual reporting sentences such as 'there is a table' but also necessarily true sentences which are 'central the community's linguistic behaviour' (Hirsch, 2005: 78), such as 'LP'. Thus, in order for the debate to be verbal, the nihilist will have to interpret 'LP' as being necessarily true in the CSO-language. Bennett's central claim is that in order for this to be possible, 'LP' would have to be *analytic* in the CSO-language (Bennett, 2009: 54). As mentioned, although Bennett states this, she does not explain *why* this would have to be the case, rather she moves immediately on to arguing why 'LP' *couldn't* be analytic in the CSO-language. Analyticity is a controversial issue, and thus we should expect some account of what Bennett is driving at when she

no table', he is not saying in his own language that 'there is no simple which is a table', because this would imply that he views his own quantifier as being semantically restricted to range over only simples. The nihilist does not regard their quantifier to be restricted in this way and so that suggestion will not work either, though Hirsch argues that the CSO speaker can charitably regard the N-language quantifier as being restricted in this way from their own perspective (Hirsch, 2005: 76-77). The question of whether the N-quantifier then is *really* restricted will itself be a verbal question – in the N-language it will be true to say 'my quantifier ranges over *everything* (not just simples)', whereas it will be true in the CSO-language to say 'the N-quantifier is semantically restricted to range over simples'. Given this, the nihilist will *not* take every assertion of the form 'there is some *x* which is *F*' as expressing 'there is some *simple* which is *F*', though the CSO can charitably regard N-assertions in this way.

¹⁶ Though the issue of *why* the equivalence holds in virtue of the sentences having the same character in the CSO-language itself raises the question of why the sentences have the same character in the CSO-language – this is looked at in the next section.

¹⁷ 'If there are simples arranged table-wise, then there is a table'.

says 'LP' has to be analytic, but no such explanation is forthcoming, and thus it then becomes difficult to assess her claim that 'LP' has to be analytic in the CSO-language if Hirsch's argument is to be successful. In this section I will fill in the following explanatory gaps then:

- i) What it would mean for 'LP' to be analytic in the CSO-language
- ii) Why it has to be the case that 'LP' is analytic in the CSO-language

Firstly then, what would it mean for 'LP' to be analytic? As a first rough gloss, we could say that in order for 'LP' to be analytic, it would need to be necessarily true, but also necessarily true in *virtue of meaning*. I assume Bennett hopes to leave this discussion at this level of abstractness but I think more detail is required on just what it is for 'LP' to be analytic.

A reasonable account of why a conditional such as 'LP' would be true in virtue of meaning would be that the antecedent necessarily entails the consequent in virtue of meaning. If the antecedent necessarily entails the consequent, the conditional itself will be necessarily true. But next, we need to understand why the antecedent would entail the consequent *in virtue of meaning*.

Consider once again the sentences 'there is a martini on the table' and 'there a mixture of gin and vermouth on the table'. Given that in the OS language the expression 'martini' is synonymous with the expression 'mixture of gin and vermouth', the truth of the former sentence will entail the truth of the latter. In this example then, it is largely uncontroversial that the former sentence, if true, entails that the latter sentence is true, and this solely in virtue of meaning. Thus, if we formed a conditional (*M*) 'If there is a martini on the table then there is a mixture of gin and vermouth on the table', then the conditional would be true in virtue of meaning. Thus let's proceed with the idea that a conditional is analytically true if the antecedent entails the consequent in virtue of meaning.

In the martini example, this was easily demonstrated, as it is uncontroversial that the expressions 'martini' and 'mixture of gin and vermouth' are synonymous in the OS language. However, that cannot be the explanation as to why 'LP' is analytic in the CSO-language. Consider that LP states 'If there are simples arranged table-wise then there is a table'. The relevant expressions to single out for apparent synonymy would be 'simples arranged table-wise' and 'table'. The problem with this, however, is that these expressions are not obviously synonymous in the CSO-language (and

making them obviously synonymous causes its own problems – see the next section for detail on this).

In what sense can ‘there are simples arranged table-wise’ entail ‘there is a table’ in virtue of meaning then? Consider Hirsch’s view that in each language a sentence is assigned a character, which assigns that sentence an intension in every context of utterance. If two sentences *S*1 and *S*2 are character-wise equivalent, then in every context of utterance where *S*1 is true, *S*2 will also be true.

Now Hirsch will say that the sentence ‘there are simples arranged table-wise’ and ‘there is a table’ are assigned the same character in the CSO-language, so they are both true in the same contexts of utterance. Thus we can say that if the CSO sentence ‘there are simples arranged table-wise’ is true, then this entails that ‘there is a table’ is also true.

What we can then say is that the sentences ‘there are simples arranged table-wise’ and ‘there is a table’ are assigned the same character in the CSO-language in virtue of the semantic rule *R* which is operative in the CSO-language (recall the discussion from chapter two, section 2.2):

(*R*): If a sentence of the form ‘there are simples arranged *x*-wise’ is true, then a sentence of the form ‘there is an *x*’ is true

If this is accepted as a semantic rule in the CSO-language, then it looks like whenever ‘there are simples arranged table-wise’ is true, then ‘there is a table’ is true – which means they will be character-wise equivalent in the CSO-language. Thus *LP* can be said to be true in virtue of meaning to the extent that the antecedent entails the consequent in virtue of meaning. Here, the antecedent entails the consequent in virtue of meaning because they are character-wise equivalent in the CSO-language in virtue of semantic rule *R* which is operative in the CSO-language. In general then we can say that a conditional is analytic if the antecedent entails the consequent, and that this entailment is explained by the antecedent and consequent having the same character in *L* owing to the rules of *L*.

So that explains what it would mean for ‘*LP*’ to be analytic in the CSO-language. Now I want to explore the reasoning for why Bennett supposes that ‘*LP*’ has to be analytic in the CSO-language.

Consider that for the debate between the nihilist and the CSO to be verbal, it has to be possible for the nihilist to interpret 'LP' as being necessarily true in the CSO-language¹⁸. If it is analytic, this gives them a nice way of doing this, but if it is not analytic, they cannot interpret it in a way without implying the falsity of their own first-order position. Let me explain.

An explanation of how 'LP' is analytic, we have seen, stems from the fact that 'there are simples arranged table-wise' and 'there is a table' are character—wise equivalent in the CSO-language. This is saying that they are intensionally equivalent in every context of utterance. Now there are two ways the sentences can be intensionally equivalent in every context of utterance.

- a) They are assigned the same character in virtue of semantic rule R
- b) They express distinct but necessarily co-obtaining facts

Now the nihilist can't say that (b) is the explanation for why the sentences are intensionally equivalent, because it implies that the nihilist is wrong in his own first-order position. For the nihilist, the only fact is that there simples arranged-table-wise; there is no additional fact which is expressed by the sentence 'there is table'. Thus he cannot account for the sentences being intensionally equivalent in every context of utterance owing to (b). This would make it look like the debate was not verbal after all and that the nihilist would be admitting that his own language was expressively deficient to the extent that he could not distinguish between certain distinct but necessarily co-obtaining facts whereas the CSO-language was capable of doing that.

Thus the nihilist should say that (a) is the explanation for how the sentences are intensionally equivalent in every context of utterance in the CSO-language. Making 'LP' analytic in this way neatly explains how the debate between the nihilist and the CSO is verbal as it turns on whether one is speaking a language where the semantic rule R is operative. The nihilist can regard the CSO-sentence 'there is a table' as being true whenever 'there are simples arranged table-wise' because the

¹⁸ We can also reconsider Bennett's preliminary worry. Though the nihilist cannot supply an undisputed falsehood that they regard as equivalent to 'there is a table', they are supposed to be able to charitably regard 'there are simples arranged table-wise' and 'there is a table' as being equivalent in the CSO-language, whilst regarding this equivalence as not holding in their own language. Thus if the dispute is verbal, this equivalence could only hold in the CSO-language if the equivalence was analytic (i.e. true in virtue of meaning). Since 'LP' is merely an object level assertion of the same equivalence, LP itself would have to be necessarily but also analytically true in the CSO-language.

semantic rule R is operative in the CSO-language. He can similarly then see that this is not a rule of his own language, and in his own language ‘there is a table’ can be regarded as being false in every context of utterance owing to character it has in his own language.

It appears that Bennett is suggesting that ‘LP’ has to be analytic for this reason then. If it is true that ‘there are simples arranged table-wise’ and ‘there is a table’ are intensionally equivalent in every context of utterance due to them expressing distinct but necessarily co-obtaining facts, this looks like it infringes on the first order debate, because it makes the nihilist language look expressively deficient. Thus ‘LP’ must be analytic¹⁹. However, Bennett claims that ‘LP’ cannot be analytic in the CSO-language. This creates a problem for Hirsch – if ‘LP’ *has* to be analytic in the CSO-language for his position to work, but Bennett can show that it couldn’t be analytic, then it looks like Hirsch’s argument fails. Let’s look at Bennett’s reasoning.

4: ‘LP’ can’t be analytic in the CSO-language

Bennett’s first reason for thinking that ‘LP’ can’t be analytic in the CSO-language can be split into two parts. The first is that if we ignore the numerical distinctness claim made by the CSO-speaker then we are not taking the debate ‘on its own terms’ (Bennett, 2009: 54). However, if we do then take into account the numerical distinctness claim, it looks as if we can define objects into existence, and Bennett supposes that analytic truths shouldn’t guarantee the existence of objects. Thus, either way, it appears that ‘LP’ being analytic creates a problem. Let’s look at the first horn of this dilemma – that if we ignore the numerical distinctness claim we are not taking the debate on its own terms.

Ignoring the numerical distinctness claim would actually explain how ‘LP’ could be analytic, and the explanation would then be much simpler. The idea is that ‘LP’ would be analytic in the CSO-language because the expressions ‘simples arranged table-wise’ and ‘table’ would be co-referring in the CSO-language, whereas they would not be in the N-language. In that case, ‘LP’ would be analytic in the same way that ‘M’ is analytic (see previous section). The problem of course is that the CSO-speaker does not take the expressions ‘simples arranged table-wise’ and ‘table’

¹⁹ There is a tension here about whether the CSO speaker will accept the claim that R is a semantic rule, and not an ontological principle of sorts, and whether *that* dispute can be given a verbal treatment. I explore such issues in the next chapter.

to be co-referring in 'LP', and to insist otherwise on this is to not take the numerical distinctness claim seriously, and we would simply end up analysing a disagreement that no one was actually having. Consider the following quotation from Bennett:

One [nihilist] says that there are only simples bearing various relations to each other, that the word 'table' is intended to refer to a composite, and thus that 'there is a table in R' is false. The other type of nihilist agrees about the ontology, but disagrees about the meaning of the predicate 'table'. He says that there are only simples bearing various relations to each other, that the word 'table' refers to simples standing in some of those relations, and thus that 'there is a table in R' is true. These two characters...under the names 'revisionary' and hermeneutic nihilist – are *both nihilists*. The dispute between them is purely verbal; it is about the semantics of English words like 'table'. However, it is not the debate that anyone is interested in. It certainly is not the debate between the nihilist and [the CSO]. (Bennett, 2009: 55)

Bennett's point here is that if we take the CSO-language to be one where the expressions 'simples arranged table-wise' and 'table' are co-referential, then the CSO-speaker is actually what Bennett calls a *hermeneutic nihilist* (Bennett, 2009: 58)²⁰. A hermeneutic nihilist would be someone who accepts the ontology of the nihilist, but also argues for a revisionary semantics for the ordinary English word 'table' (and other terms for composite objects) – thus they would say that the CSO sentence 'there is a table' is a paraphrase of the sentence 'there are simples arranged table-wise', with the latter sentence exhibiting the genuine ontological commitment. Quite clearly then, if Hirsch insisted on this he would not be analysing a disagreement between a nihilist and a common-sense ontologist; he would be, as Bennett says, analysing a disagreement between two different kinds of nihilist.

For that reason, we must be able to make sense of 'LP' being analytic in the CSO-language whilst also making sense of the idea that the expressions 'simples arranged table-wise' and 'table' are not co-referential. Bennett claims we cannot do this without thereby implying that we can define things into existence – 'surely', claims Bennett, 'an analytic claim cannot be existence-entailing in this way; surely

²⁰ This relates to the issue, raised in the previous section, about why the expression 'table' and 'simples arranged table-wise' are not synonymous in the CSO-language.

the existence of a new object cannot follow *by meaning alone*' (Bennett, 2009: 56). That being so then, we should modify 'LP' to the following:

(LP*): 'If there are simples arranged table-wise, then there is a numerically distinct table'

Just exactly why is LP* problematic? Bennett claims that it is existence entailing in a way that it shouldn't be if it was analytic, but once again she doesn't fill in the details of this claim, so I am going to examine that issue here.

Consider the conditional:

(B) 'If there is a bachelor then there is an unmarried man'

Now consider we are in a context where 'there is a bachelor next door'. The truth of this sentence, in conjunction with (B), entails that 'there is an unmarried man next door'. However, this isn't *existence entailing* because the truth of 'there is an unmarried man next door' does not express anything substantively different from 'there is a bachelor next door' in this particular context. We do not thus gain any new knowledge about the world itself, including any new knowledge regarding *what objects there are*. This is what we should expect if (B) is analytic, or true in virtue of meaning. Sentences which are just true in virtue of meaning shouldn't be able to give us new knowledge about the world itself, though they may give us new knowledge about how words are used in a particular language.

However, 'LP*' does appear to generate a problem whereby if in some context, just from knowing 'there are simples arranged table-wise' is true, 'LP*' *guarantees the existence of a numerically distinct object*. This is because individually, the sentences 'there are simples arranged table-wise' and 'there is a table' depend for their truth on the existence of distinct referents. However, from just knowing that 'there are simples arranged table-wise is true' we are guaranteed the existence of an additional object, which also then means that 'there is a table' is true. Bennett is suggesting then that 'LP*' is *existence entailing* in the following sense:

If 'LP*' and 'there are simples arranged table-wise' is true, this guarantees the existence of a referent which is not required to make 'there are simples arranged table-wise' true.

We can see that (B) is not existence entailing in this sense. For if (B) is true, and 'there is a bachelor' is true, the existence of the referent which makes 'there is an

unmarried man' is true is the *same referent* which makes 'there is a bachelor' true (relative to the same context of utterance).

What we have then is two intensionally equivalent sentences which require for their truth the existence of distinct referents. This means that the referents will have to be necessarily co-existing objects. If that wasn't the case, then there will be a possible world where one of the referents exists whilst the other one didn't, and thus the sentences wouldn't be intensionally equivalent after all. However, because semantic rule R implies that 'LP*' is necessarily true in the CSO-language, then 'there are simples arranged table-wise' and 'there is a table' are intensionally equivalent, and so there must be a necessary co-existence between simples arranged table-wise and tables. But 'LP*' is only necessarily true because we adopted a semantic rule R, thus our linguistic decision appears to have guaranteed the existence of necessarily co-existing objects. This is why Bennett remarks that 'meaning alone is enough to conjure up the existence of tables' (Bennett, 2009: 56). Thus the challenge raised by Bennett is the following – if Hirsch doesn't face up to numerical distinctness claim he is being disingenuous – analysing a debate no one is having. If he does accept the numerical distinctness claim, then 'LP*' is *existence entailing* in the problematic way that I have highlighted here.

This is a particularly worrying criticism for Hirsch because he is keen to avoid the charge that his position implies what he calls linguistic idealism, where our linguistic decisions determine what exists (Hirsch, 2002: 52). If Bennett is right about 'LP*' being existence entailing then it initially looks like Hirsch's position does imply that our linguistic decisions determine what exists.

5: Saving Hirsch from linguistic idealism

We can defend Hirsch from this particular worry, if we pay proper regard to what his views on language and reference are. Consider first of all his own explanation as to why quantifier variance doesn't imply linguistic idealism:

Consider the two sentences 'there exists something composed of Clinton's nose and the Eiffel Tower' and 'whether or not there exists something composed of Clinton's nose and the Eiffel Tower depends on our linguistic decisions'. Quantifier variance implies that the expression 'there exists something' can be interpreted in a way that makes the first sentence true or in a way that makes it false. But there

is no relevant way to interpret ‘there exist something’ that would make the second sentence true. The second sentence expresses an absurd form of linguistic idealism that is not at all implied by quantifier variance (Hirsch 2002: 52)

Let’s pick apart what Hirsch is saying here. He is claiming that

(O) ‘Whether or not there exist something composed Clinton’s nose and the Eiffel Tower depends on our linguistic decisions’

is false, whilst claiming that

(M*) Whether or not the sentence ‘there exists something composed of Clinton’s nose and the Eiffel Tower’ is true depends on our linguistic decisions²¹

is true.

That Hirsch can do this can be seen to stem from his claim that a language consists of a set of syntactic structures that are assigned characters. When a syntactic structure has been assigned a character, it is an interpreted sentence of a first-order language. Thus when speaking that particular language, we will *use* a particular interpreted sentence of that language. As such, we cannot, from within the *L* language, decide that the *L*-language sentence ‘S’ is true now rather than false. Thus consider (O). What Hirsch is saying here is that given the interpretation that ‘there exists something composed of Clinton’s nose and the Eiffel Tower’ *already* has in the M-language, we cannot just make the decision to decide that it is now false in the M-language without changing what that sentence means in the M-language. Consider that as an un-interpreted string of symbols, ‘Santa Claus exists’ can be true if we use it to mean snow is white. But given we are speaking English, we cannot just say that Santa Claus exists – that sentence already has a meaning in English and in English that sentence is false. It is in this sense that Hirsch claims that our linguistic decisions do not determine what exists. As Hirsch says, ‘once the meaning of the quantifier is fixed there is no further effect that our decisions can have on the truth-values of typical existential sentences’ (Hirsch, 2002: 53).

²¹ Here, the linguistic decision would be adopting particular semantic rules regarding the use of sentences containing quantifier expressions. Refer chapter two for more on this.

What Hirsch is saying then is that as a sentence of the M-language, ‘there exist something composed of Clinton’s Nose and the Eiffel Tower’ has a character such that it is true, and additionally (O) itself, as a sentence of the M-language has a character such that it is false in the M-language. Thus from *within* the M-language, it is not true to say that our linguistic decisions determine what exists.

However, that doesn’t mean that we cannot imagine the same syntactic structure having a different character assigned to it – this is what (M*) is attempting to capture. For we can take the string of symbols ‘snow is white’ and, knowing what it means in this language, we can nonetheless imagine that the string of symbols has a different character such that it would have different truth-conditions associated with it given that different character. This is how the A-language speaker can regard the sentence ‘there is something composed of Clinton’s nose and the Eiffel Tower’ as being a false sentence in his own language whilst being able to regard it as true in the M-language – he can see that the *same syntactic structure* has a different character assigned to it in the M-language such that it is a true sentence in the M-language. All the while it will remain false in the A-language to say that ‘our linguistic decisions determine what exists’. This then is how Hirsch attempts to deal with the problem of linguistic idealism.

However, this doesn’t yet seem to dispense with the issue we raised in the previous section. That problem was that ‘LP*’ was existence entailing in way that it shouldn’t be if it was analytic. I think this is only a problem if we are assuming a referential correspondence theory of truth, which Hirsch is not assuming. On this approach, a sentence is true only if it corresponds to some fact at a sufficiently finely grained level – the sentence will be structurally similar to the structure of the fact it is expressing. Also importantly, the semantic value of a sentence (i.e. its truth value) in a context of utterance is dependent upon the semantic values (i.e. the references) of sub-sentential expressions. Thus ‘there are simples arranged table-wise’ and ‘there is a table’ depend for their truth on different objects existing – otherwise the required correspondence for each sentence wouldn’t hold. Thus in a particular context of utterance, the truth of ‘LP*’ and ‘there are simples arranged table-wise’ guarantees the existence of the referent for the sentence ‘there is a table’. It looks like a semantic rule shouldn’t be able to guarantee the existence of referents, as the existence of a referent is an ontological matter, not a linguistic one.

However, as I mentioned, Hirsch does not accept a referential correspondence theory of truth and will certainly not agree with the claim that for a sentence to be true it has to correspond to some fact at the sort of fine-grained level that Bennett would suppose. First of all, recall Hirsch's claim that the content of a sentence is a set of possible worlds – thus the substantive content of a sentence will be, on this view, sufficiently captured at the level of the intension of the sentence. This means that the underlying structure of the sentence is not important in determining the content of the sentence – two intensionally equivalent sentences express the same substantive content, regardless of underlying sentential structure. Secondly, Hirsch argues that we can assign any character to any syntactic structure. The character of a sentence ultimately determines the truth-value of a sentence, since at every context of utterance the character will assign the sentence an intension, which again assigns the sentence a set of possible worlds. If actual world is in that set then the sentence is true. At no point in this process is the structure of the sentence itself important – the structure of the sentence does not need to correspond to the structure of the fact in order to be true. Finally consider that Hirsch i) adheres to the neo-Fregean thesis that the truth-value of a sentence is explanatorily prior to references of expressions and ii) he is committed to a deflationary theory of reference (see chapter two sections 3.1 & 4 – 4.2).

The implications of this for the present discussion are the following:

The sentences 'there are simples arranged table-wise' and 'there is a table' are intensionally equivalent in the CSO-language, and thus do not differ in their substantive content.

The character of 'there is a table' ultimately determines its truth-value in a context of utterance. Its truth-value is explanatorily prior to the reference of the expression 'table' (given neo-Fregeanism).

Finally, given the deflationary theory of reference, in any context where 'there is a table' is true in the CSO-language, the sentence 'the expression "table" refers' is also true. Thus it is not the case that the truth of 'there is a table' *depends* for its truth, in a context of utterance, on the truth of 'the expression "table" refers'. Remember that before 'there is a table' is assigned a character, it is a meaningless string of symbols. Once it is assigned a character, it will trivially follow that in any context where it is true, the sentence 'the expression "table" refers' is also true. Thus there is no explanatory priority given to the latter sentence. Given Hirsch's

views on language, true sentences about reference in a language are just trivially true and not ontologically problematic in the way that Bennett thinks they are.

Importantly, in the CSO-language, given that ‘there is a table’ is true it trivially follows that the expression ‘table’ refers, but it is not a true sentence *in the CSO-language* that ‘our linguistic decisions determine that tables exist’. As a sentence of the CSO-language, that sentence has a character such that it is false. Thus from within the CSO-language, it can be maintained that tables are not defined into existence. It can be therefore be seen that if we pay proper regard to Hirsch’s views on language, the first challenge with ‘LP*’ being analytic in the CSO-language can be met.

6: The CSO speaker does not consider ‘LP*’ to be analytic – the ‘analytic problem’

The second criticism that Bennett makes regarding ‘LP*’ being analytic is given very little attention in her article, but I think it is actually the more effective criticism, since a solution to it is not as easily forthcoming. In this section then I will expand upon the criticism that Bennett makes and show that it presents a serious challenge to Hirsch if he cannot meet it.

Bennett argues that ‘the participants in the first order debate *do not think that the relevant conditionals are analytic*’ (Bennett, 2009: 54 [emphasis added]). In this particular example, she would claim that the CSO speaker does not regard ‘LP*’ as being analytic. If ‘LP*’ is not analytic, but it is nonetheless considered by the CSO speaker to be necessarily true, then why else would they think this? As Bennett points out, the CSO speaker will believe ‘the anti-Humean claim that there are necessary connections between distinct existences’ (Bennett, 2009: 57). As such, the CSO speaker takes ‘LP*’ to be necessarily true because it describes a necessary connection between simples arranged table-wise and tables, *not merely because it is a semantic fact* about the CSO-language that the sentence ‘there are simples arranged table-wise’ and ‘there is a table’ are intensionally equivalent.

Why is Bennett’s point here important? I think it is important because it calls to attention the idea that a difference in ontological positions will manifest itself at what meta-level semantic propositions one accepts as true, thus Hirsch cannot assume what the truth-values of these semantic propositions are for the purpose of

establishing his claim that the debate between the nihilist and the CSO is merely verbal, without thereby begging the questions against at least one side in this debate. The important meta-level proposition we are concerned with here, as Bennett points out, is ‘the sentence ‘LP’ is analytic’. The CSO speaker will regard this sentence as false when applied to the object-level sentence ‘LP’ of his language – but Hirsch *needs* it to be true in order to explain how the debate between the nihilist and the CSO is merely verbal (this is only on the assumption that there is one meta-language and a number of object-level languages. I consider what would happen if we allow multiple meta-languages in the next chapter). Thus Hirsch cannot foist this upon the CSO-speaker without begging the question against the CSO-speaker.

In what follows I will expand upon and clarify what I mean by this, arguing that Hirsch’s condition of a debate being merely verbal, by only focusing on intensional equivalence between object-level sentences, has ignored an important way in which genuine ontological commitments, as opposed to superficial ones, are made using object-level sentences and that is by looking at whether those sentences themselves are considered to be analytic or synthetic. Thus Hirsch’s condition is simply not a very good condition of a debate being verbal *if* we assume that there is some important difference between an object-level sentence being necessarily true because it is analytic or because it describes a necessary feature of the world itself. Against this criticism, Hirsch’s only option is to insist that there is no real difference between understanding a necessarily true sentence as being analytic or synthetic. However, I then argue that he can only do this by adopting a verificationist argument regarding those necessary truths. I then make the point that since Hirsch has *explicitly rejected* a verificationist analysis of ontological disagreement, he can only salvage his claim that the debate between the nihilist and the CSO is verbal by retreating from his envisioned middle-ground into a thoroughgoing Carnapianism. This then, is what I call the ‘analytic problem’. I will now explain the argument in more detail.

Hirsch claims it is entirely possible for a nihilist to ‘fake it’ as a CSO-speaker, by adopting an alternative language, which I will call CSO-alt. Hirsch sets out his notion of an ‘alternative language’ in his 2009. He says that ‘in a dispute between two positions I’ll say that an “alternative language” for a given position is a language in which proponents of that position could express all of the (object-level)

propositions they believe while asserting only sentences that proponents of the other position would assert' (Hirsch 2009, 235). Thus, for the nihilist, CSO-alt is an alternative language which allows them to assert *phonetically* and *intensionally* equivalent sentences to the CSO speaker, but in such a way that they are only expressing an ontological commitment to simples²². How does this work? Well, in order for the nihilist to only express his nihilistic ontological commitments whilst speaking sentences that are phonetically equivalent to CSO sentences, the nihilist simply interprets the linking principle 'LP' as being analytic in CSO-alt. Thus the difference between CSO and CSO-alt is that the meta-level proposition 'the sentence 'LP' is analytic' is true when applied to the CSO-alt language, but false when applied to the CSO language. This means that the nihilist can assert the same object-level sentences as the CSO-speaker whilst (a) not lying about what object-level sentences he accept as true and (b) not betraying his true ontological commitments. Hirsch now observes that CSO and CSO-alt are phonetically and intensionally equivalent in terms of object-level sentences. Thus he then argues that the nihilist can charitably interpret the CSO-speaker as speaking CSO-alt, thus arguing that there is no substantive disagreement between the nihilist and the CSO.

However, this is jumping the gun since whilst Hirsch may have established that there is no difference between the nihilist language and the CSO-alt language (i.e. one could accept the nihilist ontology whilst only speaking CSO-alt), he hasn't successfully established that there is no difference between CSO and CSO-alt – he has at this point merely a) assumed there is no difference, or b) assumed that whatever this difference is cannot be the cause of a substantive disagreement. However, consider that the CSO speaker who had a genuine ontological commitment to composites would disagree with the CSO-alt speaker, who only had a commitment to simples. How would this disagreement then be drawn out, given that CSO and CSO-alt are phonetically and intensionally equivalent languages? The only way this difference can be drawn out is by *attending to the meta-level semantic propositions*. Thus, even when two languages are phonetically and intensionally equivalent at the level of object-level sentences (which may mask differences in ontological commitment), *the underlying difference in ontological*

²² The very fact that the nihilist is able to perform this kind of fakery is supposed to establish that there is no substantive disagreement between the nihilist and the CSO about how the world really is.

opinion will manifest itself in the meta-language. Consider that the CSO speaker will differentiate his genuine ontological commitments from the superficial ontological commitments of the CSO-alt speaker *precisely by insisting* that ‘the sentence ‘LP’ is analytic’ is FALSE when applied to the CSO language. Thus Hirsch cannot insist a complete identification of the CSO language with CSO-alt without begging the question against the CSO speaker.

Why is this question begging? Consider that Hirsch is supposed to be analysing how the debate between the nihilist speaker and the CSO speaker is verbal, but now he has only analysed how the debate between the nihilist speaker and the *CSO-alt speaker* is merely verbal, and attempted to identify the CSO language *as* the CSO-alt language. Since the CSO speaker will not accept this identification, precisely by claiming that the sentence ‘the sentence ‘LP’ is analytic’ is FALSE when applied to the CSO language, it looks like Hirsch cannot make this identification between the CSO language and the CSO-alt language without assuming that the CSO speaker is speaking CSO-alt, and the CSO speaker would explicitly deny this. As Bennett pointed out earlier, Hirsch must take the debate on its own terms or he risks analysing a dispute no one is actually having (Bennett 2009, 55).

Consider further, that our faking-it nihilist will be caught out if asked whether he considers ‘LP’ to be analytic or to express a necessary feature of the world itself. At this point, the nihilist can either lie or admit his conceit, either of which would demonstrate that different ontological positions will have a stake in whether ‘the sentence ‘LP’ is analytic’ is true or false. Thus the nihilist *could not* ‘fake’ being a CSO-speaker if he had to assert ‘LP’ as a synthetic truth – he *has* to understand this sentence, in his own mind, as analytic, in order to ensure he only expresses his ontological commitment to simples, whilst asserting ‘LP’. This shows, once again, that differences in ontological opinion will manifest at the level of meta-level semantic propositions, and thus Hirsch cannot merely assume the truth-values of these propositions in order to establish his merely-verbal claim without begging the question.

It is important to consider whether Hirsch envisions there being only one meta-language and a multitude of object-languages (the implications of each object-language possibly having its own attendant meta-language is considered in chapter four). If there is only one meta-language however, then it will either be true or false

in that language that ‘the sentence ‘LP’ is analytic’ when applied to the CSO-language. In order for the dispute between the nihilist and the CSO to be merely verbal, this sentence *has* to be true, otherwise the charitable interpretation of the CSO-language is unavailable to the nihilist, and thus Hirsch’s merely-verbal claim could not be established. However, if the sentence is true, then it implies that CSO is actually CSO-alt, which we find is fully compatible with a nihilistic ontology²³. Since the CSO will *not* accept this, this demonstrates that it cannot be foisted upon the CSO speaker that they are speaking CSO-alt. As Bennett argues, taking a stand on whether the linking principles are analytic or synthetic is not a neutral way of analysing the debate between the CSO and the nihilist as merely verbal, since at least one side’s ontological commitments mean that they will have a stake in whether the linking principle is analytic or not. She argues that by insisting that the linking principles are analytic, Hirsch would in fact be taking some sort of stand in the first order debate (Bennett 2009, 57). I will have more to say on this below, but *at the very least* it looks like Hirsch cannot simply assume the linking principles are analytic without thereby failing to take the debate on its own terms, foisting beliefs upon the CSO speaker that they explicitly deny, nor can he assume the linking principles are synthetic without giving up his central claim that the debate is merely verbal.

To summarise the argument thus far then, it has been shown that Hirsch attempts to identify two languages, CSO and CSO-alt, for the purpose of showing how there is no substantive disagreement between the nihilist and the CSO. This forced identification assumes that the meta-level sentence ‘the sentence ‘LP’ is analytic’ is true when applied to the CSO language, despite the fact that an important way in which the CSO speaker makes his ontological commitments clear is by denying that the meta-level sentence is true when applied to his language. What this demonstrates is that a difference in ontological commitments can actually be masked if two speakers are speaking phonetically and intensionally equivalent languages. In such case, attending to what meta-level semantic propositions each side hold as true can tease out the difference in ontological commitment. Hirsch therefore cannot really assume what the truth-values of these controversial meta-

²³²³ Recall Bennett’s example of the hermeneutic nihilist. It is entirely plausible for two varieties of nihilist to have the same ontological commitments whilst accepting different semantic claims about sentences ostensibly involving composite objects. The straight up nihilist would regard such sentences as false, whilst the hermeneutic nihilist would regard them as true but not ontologically committing.

level semantic propositions are for the purpose of establishing that the debate is verbal.

6.1: The verificationist retreat from the analytic problem

Against this, the only other option for Hirsch is to argue that there is *no difference between a necessary truth being analytic or true in virtue of some necessary feature of the world*. Hirsch would argue that whether one accepts ‘the sentence ‘LP’ is analytic’ as true or false makes no difference to what object-level sentences one asserts, nor does it make a difference to our non-linguistic behaviour, nor does it make a difference to the sorts of empirical predications we make. As a result, Hirsch can argue that if there is a difference, it is not a difference that has any discernible effect, and thus cannot be understood as a substantive disagreement. However we now need to understand the implications of arguing that there is no discernible difference between understanding a necessarily true sentence as being analytic or synthetic. As I see it, there are two conclusions one can draw from this:

I) The only discernible difference in taking a sentence to be analytic or synthetic is merely the verbal behaviour of the people involved in the debate (i.e. one side asserts the sentence is analytic, the other asserts it is synthetic). That is the only observational outcome. There is no difference in what object-level sentences one regards as true, and no difference in the empirical predictions one makes by taking a necessarily true sentence to be analytic rather than synthetic, and vice versa. Thus one could argue that since there is no discernible difference beyond the initial difference in verbal behaviour, then there is no *actual difference*. I think this conclusion is unattractive for Hirsch since if we get rid of a neat divide between analytic and synthetic sentences, we also get rid of a neat divide between substantive and merely verbal debates.

Going down this route is not attractive to Hirsch for a number of reasons. Firstly, if there is no neat dividing line between substantive and verbal debates, then Hirsch cannot use this to then definitively establish that the debates in physical object ontology are merely verbal. However, it is quite obvious that Hirsch wants to maintain that there is some intelligible distinction between substantive disputes and merely verbal disputes – remember that he envisions some ontological disputes as being substantive (such as the Platonism/nominalism debate – see Hirsch 2009,

256-6, for example). Thus getting rid of a principled distinction between analytic and synthetic truths will not allow Hirsch to maintain the position he has carved out for himself. Equally, if Hirsch is hoping to be a defender of common-sense (as mentioned as one of his central aims – see section 7 of the previous chapter), then it looks like he will have to maintain the analytic/synthetic distinction since it seems to me that an important facet of common sense is an understanding that some debates are definitively factual and some are definitively verbal.

Furthermore if we deny that there is any actual difference between a sentence's being analytic or synthetic based on the idea that there are no observational differences (beyond the linguistic behaviour involved in asserting or denying the meta-level semantic propositions themselves), then it looks like Hirsch is making some kind of retreat into verificationism. There are a number of different ways of understanding verificationism, but at its most basic if we understand it as the claim that the meaning of a sentence is its method of verification, then there are no means of definitely establishing whether the sentence 'the sentence 'LP' is analytic' is true or false (since the outcomes of it being true or false are the same). Against this, it could be argued that the verificationist principle is only intended to apply to object-level sentences. However, given that certain necessarily true object-level sentences are supposed to be about the world, then there would be no finite verification conditions for necessarily true sentences, meaning the linking principles of the CSO-language would either being meaningless or taken to be mere linguistic conventions of the CSO-language. This creates several problems for Hirsch. First, it once again runs into the problem that the CSO speaker will deny that their linking principles are mere linguistic conventions. Hirsch can only insist that they are by retreating into verificationist justifications, thus abandoning his avowed rejection of a verificationist analysis of ontological debate (see section 7 of the last chapter).

This also causes a further issue. Consider that the whole purpose of CSO-alt is that the nihilist is able to express the object-level propositions consonant with their nihilistic ontology whilst asserting sentences that are phonetically and intensionally equivalent to the CSO speaker. However, by retreating into verificationism, such that all object-level necessary truths are taken to be mere linguistic conventions, then there is no sense in which the nihilist, now speaking CSO-alt, is actually still a nihilist. All he has done, according to the verificationist,

is revise his own (originally nihilistic) linguistic conventions. The idea that the nihilist can 'fake it' as a CSO-alt speaker then, whilst still being secretly a nihilist, makes little sense when his own nihilistic commitments are nothing more than linguistic conventions.

To summarise the first problem then, if Hirsch goes after the analytic/synthetic distinction in this way, he retreats from several commitments that are seemingly central to his position. The first is that there is now no clear cut distinction between *any* debate being merely verbal or factual. This goes against Hirsch's original claim that the debates in physical object ontology are merely verbal, since this claim cannot be definitively established. Secondly, this means other debates that Hirsch considers to be definitely substantive are no longer definitively substantive. Thirdly, taking the debate in this direction can be seen to be a retreat into verificationism, or at least draws on verificationist-style arguments. This is problematic because Hirsch explicitly denies that he is a verificationist. It also further complicates Hirsch's claim that the nihilist could speak CSO-alt whilst still being a covert nihilist. This no longer makes sense on a verificationist account since there is nothing more to being a nihilist than accepting certain linguistic conventions. Thus by switching to being a CSO-alt speaker, the nihilist couldn't be called a nihilist in any meaningful sense.

II) A further option available to Hirsch is to argue that whilst there may be a difference between 'LP' being analytic or synthetic, there is simply no way of establishing this. However, this doesn't really prove that the debate between the nihilist and the CSO is merely verbal. It just proves that it is hard or possibly impossible to resolve. However, again Hirsch explicitly rejects this analysis of the debate (Hirsch 2009, 231).

Against Hirsch's arguments here, those sympathetic to a substantive approach to physical object ontology will argue that whilst there is no difference between speaking CSO and CSO-alt that will satisfy the verificationist, an appeal to well-known philosophical devices will elucidate the differences. Consider that if 'LP' is analytic, it does not require the world to be a particular way in order for it to be true – i.e. the argument can be made that analytic truths are those that do not require a truthmaker. However, if 'LP' is not analytic, it *does* require the world to be a particular way. We can thus make an appeal to the idea that if 'LP' is true and *not*

analytic, there will be some truthmaker that makes it true. However, if it is true but merely analytic, there is no requirement for there to be a truthmaker. Although such appeals will not satisfy a verificationist, Hirsch construes his position as being anti-verificationist²⁴.

Thus ultimately the analytic problem reveals that Hirsch cannot maintain that the debates are merely verbal without thereby giving up an important aspect of his position. This fits into my more general point: Although a thoroughgoing Carnapian approach, with its verificationist/anti-realist commitments may yet be a viable alternative to substantive approaches to ontology, Hirsch's modified Carnapianism is shown up to be an untenable position due to his avowed rejection of a verificationist analysis of ontological debate.

7: Summary of chapter three

In this chapter the main point that Bennett raises against Hirsch is that in order for his position to be verbal, the sentence 'LP*' has to be analytic in the CSO-language. First of all, I explained what it means for 'LP*' to be analytic in the CSO-language – a point Bennett herself does not discuss. I argued that 'LP*' could be regarded as analytic if we consider that 'there are simples arranged table-wise' entails 'there is a table' in virtue of meaning, which I then further analysed as meaning that it was true that the former sentence entails the latter in the CSO-language owing to the semantic rule R being operative in the CSO-language. Thus 'LP*' could be seen to be necessarily true as owing to nothing more than the semantic rules operative in the CSO-language. I then explained, again filling in the explanatory gap left by Bennett, *why* 'LP*' has to be analytically true in the CSO-language. I explained that the nihilist needs to be able to charitably interpret 'there are simples arranged table-wise' and 'there is a table' as being intensionally equivalent in the CSO-language. They could only do this if the sentences were intensionally equivalent in virtue of meaning, and *not* in virtue of the fact that they express distinct but necessarily co-obtaining facts.

Bennett then argued that 'LP*' *can't* be analytic in the CSO-language, firstly because it would imply that we could define objects into existence. I demonstrated

²⁴ I am not arguing that truthmaker theory is the only way of elucidating the difference between someone who holds a sentence to be analytic versus someone who regards it as synthetic. However, it does provide a way of showing the difference between the CSO and CSO-alt speaker that can only really be barred by appealing to verificationist style-arguments.

that this particular issue could be met by Hirsch if we pay proper regard to his views on language. On that view, a sentence's truth-value in a context is determined by its character, and thus there is no explanatory priority given to the references of sub-sentential expressions in determining whether a sentence is true or not.

However, Bennett also remarks that the CSO speaker does not regard 'LP*' as analytic. This, I argued, was a far more interesting point and developed this into an argument I called the analytic problem. Essentially, this is the idea that differences in ontological opinion will manifest at the level of the meta-language. Thus Hirsch cannot assume the truth value of those sentences for the purpose of establishing that the debate is verbal without begging the question. This calls into question how suitable Hirsch's initial condition for a debate to be verbal is. Against, this I argued that Hirsch could only solve the analytic problem by retreating into a thoroughgoing Carnapianism. This supports my central contention that Hirsch's middle-ground position is untenable.

Chapter four: A potential solution to the analytic problem

1: Introduction

In the previous chapter I looked at Bennett's criticism from Hirsch, and from that I developed what I called the *analytic problem*. The analytic problem stems from the idea that differences in ontological commitment will manifest themselves not only in object-level sentences but also in a disagreement over meta-level semantic propositions. By focusing on object-level sentences, Hirsch's criterion for a debate being verbal ignores the ways in which a debate about what objects exist will carry over into the meta-language regarding claims over whether relevant object-level sentences are analytic or synthetic. As such, I argued in the previous chapter that Hirsch cannot simply assume the truth-values of these controversial meta-level claims without begging the question against at least one side in the debate. This is because it was shown that there could be two languages that were phonetically and intensionally equivalent, such that the underlying difference in ontological commitments between the languages was masked. It was only by looking at what meta-level semantic propositions each side adhered to that the difference in ontological commitment could be teased out. Thus Hirsch could not assume the truth-values of the relevant meta-level semantic propositions in order to make his merely-verbal claim, since this would then presume that one side in the ontological debate was actually speaking a language that they explicitly denied they were speaking.

The relevant semantic proposition in question, in the debate between the nihilist and the CSO is (AS) 'the sentence 'LP*' is analytic'. As it was shown in the previous chapter, the nihilist needs this claim to be true if they are to be able to charitably interpret the CSO speaker. Equally, the nihilist needed this sentence to be true if they were to 'fake it' as a CSO speaker whilst still only expressing the object-level propositions consistent with a nihilist ontology (which Hirsch claims is entirely possible). However, the CSO-speaker explicitly denies that AS is true when applied to the sentence 'LP*' of the CSO language. Thus Hirsch needs AS to be true in order to analyse the debate as verbal, however this seems to assume that the CSO speaker has no stake in the truth of AS. The previous chapter showed that this was not the case, and thus the analytic problem poses a genuine problem for Hirsch.

At the end of the previous chapter, I discussed a number of potential ways Hirsch might deflate the analytic problem. However, these ultimately amounted to either rejecting a clear-cut analytic/synthetic distinction for necessarily true sentences, or to making verificationist style arguments to render the commitment to a necessarily true sentence as being analytic or synthetic meaningless or otherwise inconsequential (and at the very least not substantive). Both these ways of deflating the analytic problem are certainly viable, but the problem for Hirsch is that they involve a *retreat from his position as he presents it*. Remember it is the central claim of this thesis that whilst thoroughgoing Carnapian treatments of ontological debates might be viable, Hirsch explicitly sets himself up as *not* being a thoroughgoing Carnapian, and attempts to adopt a modified light version of Carnapianism. In the previous chapter I argued that the ways out of the analytic problem discussed there involve too severe a retreat for Hirsch, showing that his modified position is untenable.

In this chapter, I want to explore to further ways Hirsch might seek to deflate the analytic problem. The first is to argue that Hirsch's position need not make a commitment to 'AS' – Hirsch can remain somewhat agnostic about this whilst still retaining that the debate between the CSO and nihilist is merely verbal. My argument against this strategy is that Hirsch simply cannot avoid the issue in this way when it is demonstrated that ontological positions will have a stake in the truth value of 'AS'. The second option is to argue that the debate over the truth value of 'AS' is *itself merely another verbal debate*, such that the answer to that question depends on what language one is speaking.

This second option is the more plausible route for Hirsch, and I argue it is more in line with similar arguments he has made elsewhere. For example, it will be recalled from chapter two that a debate over whether an object, *a*, exists, carries over into a debate over whether the expression 'a' genuinely refers. Hirsch's strategy is to argue that if the former is merely verbal then so is the latter. In that chapter it was seen that in order for this to be the case, Hirsch needed to be committed to a deflationary theory of reference. In this chapter, I show that in order for the debate over the truth-value of 'AS' to be merely verbal, Hirsch will have to argue that whilst there is a *nominal distinction* between necessarily true object-level sentences that are analytic and ones that are synthetic, this distinction is only *illusory*. In reality, all necessary truths will turn out to be analytic. In one respect, this result

is good for Hirsch because it allows him to offer a further diagnosis of why the nihilist and the CSO mistakenly think their dispute is substantive – they think there is a genuine difference between necessarily true object-level sentences that are synthetic and those that are analytic.

However, in adopting the view that all necessary truths are in-fact analytic, Hirsch would be creating a number of issues for himself. There are two which are relatively minor (though definitely problematic in their own right) – these issues will be explored in the relevant section of this chapter. The major problem with making all necessary truths analytic is that it pushes Hirsch toward a thoroughgoing Carnapianism. I will argue that Hirsch must give up his claim to be a robust realist (and by proxy his commitment to being a defender of common sense) and he must also give up his claim that his merely-verbal claim only applies to debates in physical object ontology. Thus this new solution to the analytic problem once again demonstrates that Hirsch must retreat from his modified light Carnapianism into a more thoroughgoing variety, once again demonstrating his position is untenable.

2: The first solution: Do not commit to ‘AS’ being analytic or synthetic

In this section I will explore the first potential solution to analytic problem. It is quite possible that Hirsch could simply deny that the debate over whether ‘AS’ is analytic or not is simply inconsequential to establishing that the ontological dispute between the nihilist and CSO is merely verbal. Rather than making an explicit commitment to the truth-value of ‘AS’, rather Hirsch could simply argue that as long as it is demonstrable that the nihilist language and the CSO language are intensionally equivalent (in terms of object-level sentences), then it is fine for one side to *usefully regard* the ontological principles of the other side as being merely semantic rules (regardless of what the other side actually thought of the status of the relevant sentences in their own language).

For Hirsch, what is of central importance in establishing that the debate between the nihilist and the CSO is verbal is that each side can ‘plausibly interpret the other side as speaking a language in which the latter’s asserted sentences are true’ (Hirsch 2009: 231). This would of course include sentences which one side took to be necessarily true. If we are the N-speaker then we need some plausible explanation of the CSO-speaker’s linguistic behaviour, i.e., the fact that they would be willing to assent to the sentence ‘LP*’ being true in any context of utterance. As long as the nihilist can come up with this plausible explanation, they can use it to

charitably regard the CSO as speaking the truth in their own language, even if the CSO seemingly disagrees with the explanation that the N-speaker has given.

Consider what Hirsch says in the following quotation about a mereological essentialist trying to interpret sentences that a four-dimensionalist would regard as necessarily true:

It seems obvious that the principles of mereological sums and temporal parts are in some sense central to the DL-community's *linguistic behaviour*. From our own mereological essentialist perspective we can *usefully regard* those principles as working in effect as semantic rules that generate truth conditions for the disputed sentences... A disputed sentence is true in DL-English if it follows from the undisputed facts in conjunction with the two principles. An immediate consequence is that the principles themselves are (a priori) true with respect to every possible situation, which is of course what the DL-community says. (Hirsch, 2005: 78 [emphasis added])

In the above passage, Hirsch is suggesting that what is important from the ME perspective is that we find *some way* of charitably interpreting sentences that are central to the 4D-speaker's linguistic behaviour. But what does it mean to say that a sentence is central to a community's linguistic behaviour? The ontological principles that the 4D-speaker accepts, for example, that ordinary objects are composed out of temporal parts, are central to the 4D-speaker's linguistic behaviour because they seem to dictate the syntactic form that perceptual reports²⁵ take in the 4D-language. For example, we observe that in every context of utterance, a 4D-speaker will assent to 'objects are composed out of temporal parts' and we see that in a specific context of utterance, the 4D-speaker says 'the temporal part of the tree at t_1 is brown'. Thus, in a situation where we would say 'the tree is brown', the intensionally equivalent form that a 4D sentence must take in the same context appears to be determined by the fact that they accept the sentence 'objects are composed out of temporal parts' as being necessarily true. It is in that sense that the sentence is 'central to the DL-community's linguistic behaviour'. From the ME perspective then, we need to explain *why* the sentence

²⁵ The sense in which I use 'perceptual report' here is the same sense in which Hirsch uses it in his 2005. See chapter two.

‘objects are composed out of temporal parts’ is central to the 4D community’s linguistic behaviour in this way. The competing explanations are:

- i) The sentence is a necessarily true ontological principle
- ii) The sentence is a semantic rule

From our own perspective (i) is clearly not a good explanation, because in our language ‘objects are composed out of temporal parts’ is false. Thus a better explanation of the sentence’s centrality in the 4D-language is that it is a semantic rule of the 4D-language. By regarding it as such, we make good sense of *why* the sentence is always regarded as true, and we also make sense of the form that perceptual reports take in the 4D-language. That is, from our own perspective, we can plausibly interpret ‘objects are composed out of temporal parts’ as a semantic rule of the 4D-language, because doing so is in accordance with the fact that they regard the sentence as necessarily true, and that it seems to dictate the form that other sentences in the language take. In other words, the explanation fits the observation.

Of central importance here is the notion of a sentence’s character, which is a function from context of utterance to a proposition, which is for Hirsch a set of possible worlds. The character of a necessarily true sentence is such that it will deliver, in any context of utterance, the proposition which is the set of all possible worlds – i.e. the sentence will be true in every situation. Whether or not a sentence is a semantic rule or an ontological principle *makes no difference* to the *character* of the sentence. Thus, although the 4D-speaker considers ‘objects are composed of temporal parts’ as being an ontological principle, we can usefully regard it as a semantic rule because, as I said above, it fits the observation, but more importantly this also makes no difference to the character of the sentence.

Furthermore, recall that for Hirsch a charitable interpretation of a linguistic community’s assertion must rationally reflect their causal relation to their environment (particularly their perception of that environment) (Hirsch, 2005: 73). As such, we should interpret a sentence uttered in a context in a way that makes sense of how the speaker is causally related to the environment²⁶. Now, consider a sentence ‘S’ that the speaker considers necessarily true. Usefully regarding that sentence as a semantic rule instead of an ontological principle is not a better, but nor is it a worse interpretation of the sentence when we are considering how a

²⁶ For more details on this, refer back to the sections 6.1-6.3 in chapter two.

speaker who holds that sentence to be necessarily true is causally related to their environment. This is because if a speaker originally considers 'S' to be an ontological principle, but later switches to considering 'S' to be a semantic rule, this in itself will change no aspect of their non-verbal behaviour or how their assertions rationally reflect their causal relation to the environment. They will still assert it as being true in every context of utterance, and it will not change what perceptual reports they regard as true or false, and it will not change what empirical predictions a speaker is likely to make. For example, imagine someone who says 'if I go home tonight, there will be a temporal part of the tree in the garden²⁷ which is brown'. This is an empirical prediction they are making - they expect to see a temporal part of the tree in the garden which is brown later on in the day. However, let's imagine this person speaks like this first because they believe the sentence 'objects are composed of temporal parts' is an ontological principle, but later during the day he then considers the sentence to be a semantic rule. This will in no way affect in what contexts he regards 'objects are composed of temporal parts' as true, nor will it cause him to revise his empirical prediction, expressed by the sentence 'if I go home tonight, there will a temporal part of the tree in the garden which is brown'. Thus either way, interpreting a necessarily true sentence as a semantic rule rather than an ontological principle makes no difference regarding how its assertion by a speaker reflects that speaker's causal relationship to their environment.

Thus we can usefully regard a sentence 'S', which is considered to be an ontological principle, as a semantic rule for three reasons:

- 1) It explains how perceptual reports made in that language have the specific form that they do.
- 2) It does not change in what contexts the speaker would regard the sentence as true because necessary truths have the same character whether or not they are semantic rules or ontological principles.
- 3) Regarding 'S' as a semantic rule does not change how its assertion by a speaker reflects how that speaker is causally related to their environment as a speaker can switch from believing 'S' is an ontological principle to believing

²⁷ Strictly speaking of course, the home and the garden will also be composed of temporal parts.

it is a semantic rule without affecting their non-linguistic behaviour, the perceptual reports they make, and the empirical predictions they make.

Now then, from the nihilist perspective we can *usefully regard* 'LP*' as being analytic in the CSO-language. Doing so is in line with the three reasons above:

- 1) 'LP*' being analytic in the CSO-language explains how perceptual reports in that language have the specific form that they do.
- 2) It does not change in what contexts the speaker would regard the sentence as true because necessary truths have the same character whether or not they are analytic or ontological principles.
- 3) Regarding 'LP*' as analytic does not change how its assertion by a speaker reflects how that speaker is causally related to their environment as a speaker can switch from believing 'LP*' is an ontological principle to believing it is analytic without affecting their non-linguistic behaviour, the perceptual reports they make, and the empirical predictions they make²⁸.

With this reasoning in place then, it looks like Hirsch doesn't need to explicitly make a stand on whether the linking principles in the CSO language are analytic or not. Rather, all he needs to do is show that it is entirely plausible for the nihilist to usefully regard the principles as such. Since he can do that, he does not need to explicitly state either way whether 'AS' is true or false.

However, after considering what was discussed in the previous chapter, it should be seen that Hirsch can't simply side step the question of the truth-value of 'AS' in this way. This is because we can no longer assume that meta-level semantic propositions are inconsequential to the debate between the nihilist and the CSO, since different ontological commitments will manifest themselves in a dispute over the truth value of sentences like 'AS'. Quite simply then, Hirsch cannot argue that meta-level semantic propositions are inconsequential when those very propositions are seen to be at stake in the ontological debate.

²⁸ It will be noted here that I have switched from saying 'semantic rule' to 'analytic'. In the section on Tractarian analyticity I will explain the idea that whether a sentence is a semantic rule or analytic amounts to the same thing – analytically true sentences are just necessarily true rules which constrain the form that sentences describing our experiences can take.

3: The second solution: The dispute over 'AS' is itself merely verbal

The more viable solution for Hirsch is to argue that the dispute over the truth-value of 'AS' itself can be seen to be a further part of the original verbal debate. How can this be done?

It will be recalled that in the previous chapter, the truth value of 'AS' only became problematic if we assumed that there was *one* meta-language and a number of different object-level languages. However, it could be argued that just as each side is making object-level assertions in their own object language, they are also making meta-linguistic assertions in their own meta-language. The basic idea then is that just as there is a CSO object-language, there is an attendant CSO meta-language and in that meta-language the meta-linguistic claims of the CSO speaker come out true. Equally there will be a nihilist object-language and nihilist meta-language. The idea is that the nihilist can make claims in the nihilist meta-language *about* the CSO object-level language, and vice versa. For example, if we recall from the previous chapter, in order for our nihilist to 'fake it' as a CSO speaker whilst only expressing their commitment to simples, they need to understand the linking principles in the CSO language as being analytic. Thus the nihilist can maintain, *in their own meta-language*, that the sentence 'AS' ('The sentence 'LP*' is analytic') is true. This way the nihilist can assert 'LP*' in the company of CSO-speakers, on the understanding that according to what is true in his own meta-language, 'LP*' is an analytic truth. The nihilist can thus assert the object-level sentence without thereby committing to the existence of composite objects. Equally, 'AS' can be regarded as *false* in the CSO meta-language. The idea then is that the speaker's attitude toward necessarily true-object level sentences in their own object-language comes out right in the speaker's own meta-language.

Ultimately, the idea is that the truth-value of 'AS' depends on what meta-language one has adopted, and thus the debate over the truth-value of 'AS' is *itself merely verbal*.

This is quite a neat solution for Hirsch because it nicely deals with the apparent disagreement over the truth-values of meta-level semantic propositions. He can now say that if the debate over whether composite objects exist carries over into a disagreement over 'AS' and other such sentences, then if the former debate is merely verbal then so is the latter (just as he argued that debates over whether a term 'a' genuinely refers are as verbal as the debate over whether the object *a*

exists'. It also deals nicely with a problem raised in the previous chapter. According to Hirsch, it is entirely possible for the nihilist to fake it as a CSO speaker. How they do this is now explainable as the nihilist adopting the object-language of the CSO speaker whilst retaining the nihilist meta-language. The idea then is that whilst the nihilist and the CSO speaker converge on what object-level sentences they assert, they will still diverge on what meta-level sentences they assert – but that is now to be seen as nothing other than an additional linguistic decision and nothing more (i.e. the initial linguistic decision will be to switch from speaking the nihilist object-language to speaking the CSO object-language. It is then merely a further linguistic decision to switch from speaking the nihilist meta-language to the CSO meta-language).

I will now argue that this is a plausible solution for Hirsch. In the previous section I argued that whilst regarding the sentence 'LP*' as analytic rather than being an ontological principle does not change what perceptual reports one makes, what empirical predictions one makes, nor does it change any non-linguistic behaviour, it will have a knock-on effect regarding what *other* sentences are regarded as true – thus accepting 'LP*' as analytic has an effect on linguistic behaviour, but not non-linguistic behaviour. Specifically, if one thinks 'LP*' is analytic, they will assent to 'AS' and if one thinks 'LP*' is *not* analytic, they will *not* assent to 'AS'. 'AS' is an example of what I will call an *attitude-sentence* because it reports a certain attitude that a speaker has toward a necessarily truly sentence, namely, whether they regard it as being analytic or not. Considered in schematic form an attitude-sentence would look like this:

(AS): 'The sentence ϕ is analytic'

The idea is that we can substitute ϕ for an object-level sentence considered to be necessarily true in a language, and this will have an effect on whether AS, as a whole, is considered true or false in the meta-language. For example, if we replace ' ϕ ' with 'all bachelors are unmarried males', the CSO-speaker will regard AS as true. However, if we replace ' ϕ ' with 'LP*', the CSO-speaker will regard AS as false.

I'm going to suggest now then that whether or not 'AS' is true or false *depends on what meta-language one is speaking*, and that it is possible to charitably interpret an AS sentence as being true in another meta-language whilst holding that it is

false in your own meta-language²⁹, and thus will show how one can argue that the question of whether ‘LP*’ is analytic is itself *merely a verbal debate*.

Consider that the CSO will regard (AS2) ‘the sentence ‘LP*’ is synthetic’ as true in their own meta-language. If we consider in what contexts the CSO will regard ‘AS2’ as true, we will find that they regard it as true in every context of utterance. Hence in the CSO meta-language it appears as if ‘AS2’ *itself* is held to be necessarily true. Thus this particular attitude sentence itself can be usefully regarded, from the nihilist perspective, as being true in the CSO meta-language *for all the same reasons* that we originally charitably regarded ‘LP*’ itself as being a semantic rule in the CSO-language. Consider that if one switches from thinking ‘AS2’ is true to later thinking it is false (without changing their view that ‘LP*’ itself is necessarily true), then this does not in any way affect the *character* of the ‘LP*’ – it is still regarded as true in the same contexts, and it does not change how the speaker is causally related to the environment. Therefore, the nihilist can charitably interpret the sentence ‘AS2’ as being necessarily true in the CSO meta-language in just the same way as they would charitably interpret ‘LP*’ as being necessarily true in the CSO object-language – both sentences are now seen to be central to the linguistic behaviour of the CSO speaker, and thus both can be regarded as being ‘semantic rules’ or true in virtue of nothing other than linguistic decision in the CSO meta-language and object language respectively. In doing this, the nihilist can charitably regard the CSO-speaker’s attitude toward ‘LP*’ as coming out right in the CSO meta-language. Equally, the nihilist can now understand their own attitude toward ‘LP*’ as coming out right in the nihilist meta-language.

Thus from the standpoint of a particular meta-language, we can maintain a distinction between necessarily true object-level sentences which are regarded as analytic and necessarily true object-level sentences which are considered to be synthetic. Thus the CSO speaker, from the standpoint of their own meta-language can maintain a distinction in the CSO object-language between ‘LP*’, which is taken to be synthetic, and sentences such as ‘all bachelors are unmarried men’, which is taken to be analytic. Ultimately then, it will be seen that the dispute over the truth-value of ‘AS’ is itself merely another manifestation of what is ultimately a verbal debate, since it is possible for the nihilist to regard ‘AS’ as being true in their

²⁹ Just as it would be fine for the nihilist to regard the object-level sentence ‘there is a table’ as false in their own language whilst charitably interpreting it as true in the CSO object language.

own meta-language (this they need to do for the purpose of charitably regarding 'LP*' as coming out true in the CSO object-language), whilst charitably regarding 'AS2' as being true in the CSO meta-language.

4: Problems with explaining the truth-values of attitude sentences

We have so far seen that a potential solution to the analytic problem is that the relevant attitude sentences which ascribe or deny analyticity to necessarily true sentences can themselves be charitably regarded as true in the relevant meta-languages. Thus each side can retain a *prima facie* correct analytic/synthetic distinction between object-level necessarily true sentences of their own object-language, which is in line with their ontological commitments.

The implication of this solution is that nothing appears to determine or explain the analyticity or the syntheticity of necessarily true object-level sentences beyond whether a particular attitude sentence about that sentence is true. Thus nothing appears to determine or explain that 'LP*' is synthetic in the CSO-language beyond the fact that 'AS2' is true in the CSO meta-language. We might normally suppose that analyticity or syntheticity are semantic properties of sentences that are determined by the semantic properties of the parts of the sentence. Thus the analyticity of the sentence 'all bachelors are unmarried men' is explained bottom-up from the meanings of 'bachelor', 'unmarried' and 'men', and the arrangement of these terms. Thus we might suppose that it can't simply be a result of stipulation in the meta-language that 'all bachelors are unmarried men' is analytic – whether the sentence is analytic or not will depend on the semantic facts about the sub-sentential expressions in the sentence and their arrangement. This suggests that there is some objective criterion or condition that a sentence must fulfil if it is to have the property of being analytic or synthetic.

We might, for example, suggest that a sentence is analytic if it states that two sub-sentential expressions within the sentence are synonymous (such as "all bachelors are unmarried men"). Call such a criterion for now the *strict analytic criterion*, or SAC. If we argue then that it is a condition of a sentence being analytic that it fulfils SAC, then evidently 'LP*' is not analytic because 'table' and 'simples arranged table-wise' are not regarded as synonymous expressions. However, the problem with this is that if SAC is a criterion for a sentence being analytic, then whether or not 'LP*' is analytic can't be verbal – the dispute will turn on whether the sub-

sentential expressions of 'LP*' are synonymous in the CSO-language. They are not, and so the nihilist could not maintain 'AS' in their own meta-language.

However, if we consider Hirsch's neo-Fregean position, such that the truth-values of sentences are explanatorily prior the semantic properties of sub-sentential properties, there is nothing wrong with supposing that there is no substantive explanation of an object-level sentence 'S' being analytic or synthetic beyond the truth of some sentence in the meta-language which asserts that 'S' is analytic/synthetic. We should remind ourselves here of the parallels to the debate over reference. Quite simply if '*a* exists' is true, then the sentence 'the expression '*a*' refers' is also true, and vice versa – neither of these sentences is explanatorily prior to the other. There is no special substantive relation that '*a*' bears to extra-linguistic reality that explains the truth of 'the expression '*a*' refers'. Equally then, we might argue that there is no special relation between the sub-sentential expressions of 'LP*', or no special relation between the sentence as a whole and its relation to reality, that explains the truth of 'AS2'³⁰ in the CSO meta-language.

We now need to turn to the following issue: What then determines the truth-values of the relevant attitude sentences in the meta-language if there is no genuine property or relation pertaining to the necessarily true object-level sentence that explains its analyticity/syntheticity? For we observe that the character of 'LP*' itself does not determine that it is analytic or synthetic – the character of 'LP*' will be the same in either case – the character merely determines 'LP*' as true in every context of utterance. Thus there is nothing about 'LP*' itself that explains why it is regarded as synthetic in the CSO object-language, whilst another necessarily true sentence such as 'all bachelors are married men' is regarded as analytic. Thus it would appear that all that determines why 'AS2' is true in the CSO-language is the *character* of 'AS2' itself, and we find that 'AS2' *itself* has a character in the CSO-language such that it is necessarily true.

The question now is what *determines the character* of 'AS2'? Since the question of whether 'AS2' itself is true or false is now considered a merely verbal issue, and 'AS2' will be true or false depending on what language one is speaking, then whether or not 'AS2' is true or false must itself be determined by nothing other than a *linguistic decision*.

³⁰ The sentence 'LP*' is synthetic' – remember the CSO does not want 'LP*' to be analytic.

I now want to clarify what I mean exactly by a linguistic decision here before explaining why this potentially generates an issue for Hirsch, as it begins to look like *all* a priori necessary truths are true in virtue of nothing more than linguistic decisions, and in the next section I will argue that this is tantamount to saying that all necessary truths are analytic (I will explain in the next section the specific sense of analyticity) This outcome poses a number of issues for Hirsch which I will explain in section 6. In chapter two, section 4 I argued that for Hirsch a language is just a set of characters assigned to a set of syntactic structures to give us meaningful, interpreted sentences. The character of a sentence determines in what contexts of utterance a sentence is true or false, and thus character ultimately determines whether a sentence is true or false. Once the character of a sentence is fixed, we cannot decide in what contexts the sentence is true or false – how the world is will determine whether or not, in a particular context, a sentence is true or false. In that case we might say that a character of a sentence indirectly determines its truth value – for it determines in what contexts the sentence is true, but of course the sentence may or may not be asserted in those contexts. However, necessarily true sentences are true in every context of utterance, and so we can know beforehand, in any context, that the sentence will be true. Thus the character of a necessarily true sentence *directly* determines that it is true, irrespective of context.

The question now then is what is it that determines that a certain sentence has a character such that it is necessarily true? The answer can only be *that the linguistic practices of the community determine it to be necessarily true*, i.e. the way the linguistic community uses the sentence explains why it has the character that it does have, and thus the way the sentence is used in the linguistic community determines that it is necessarily true.

Let's now briefly recap what I mean by linguistic decision, which I initially discussed in chapter two, section 4.1. For Hirsch, we can see that there are two kinds of linguistic decision. The first is the decision to speak a certain language – for example we might decide to switch from speaking the CSO-language to the N-language. In this kind of linguistic decision, we move from speaking one interpreted language to speaking another language which is itself already interpreted – the characters of the sentences in the new language are already fixed. Thus I cannot go from speaking the CSO-language to speaking the N-language and claim that the sentence 'there is a table' is now true in the N-language – that sentence already has

a character in the N-language. In one regard then, when Hirsch says whether ‘there is a table’ is true depends on our linguistic decision, he can mean that it is up to us whether we speak the CSO-language or the N-language.

However, there is another sense in which we can use ‘linguistic decision’ – this refers to the shared linguistic behaviours of a linguistic community that determine the characters of sentences of their language in the first instance. How a linguistic community *uses* a certain sentence originally will determine the character of that sentence for the linguistic community. Call this then the meaning determining behaviour of a linguistic community. The meaning determining behaviour of a community will determine that some sentences are always true regardless of how the world is or in what context the sentence is asserted – thus they will be determined as being necessarily true.

This has to be the correct meta-semantic explanation for how sentences get their character according to Hirsch, as it is what allows him to say that the debate between the CSO and the nihilist is merely verbal. If his explanation is that they are speaking different languages where there is a permutation of characters of the same syntactic structures, such that in the CSO-language ‘there is a table’ is true in some contexts but in the N-language this sentence is necessarily false, then there can’t be any substantive reason why the characters are permuted in this way, otherwise the debate would be substantive – the only explanation then is that the characters of sentences are determined by how the sentences are *used* in a linguistic community – i.e. they are determined via a linguistic decision of the second kind. This has the consequence that necessary truths are determined as such *purely by linguistic decision*, since the context in which a necessary truth is uttered cannot affect its truth value. What consequences does this have for Hirsch and the present proposed defence against the analytic problem?

5: Tractarian-analyticity

In the previous section it was found that the dispute over whether ‘LP*’ was analytic could be determined as merely verbal by seeing that ‘AS’ could have different truth values in different languages. It was then enquired as to what determined the truth value of ‘AS’ in the respective languages, and it was argued that the truth value was not determined by any property of ‘LP*’ itself. Thus the only thing that could determine the truth value of ‘AS’ was its character. It was then argued that being necessarily true, the character of ‘AS’ determines that it is

true in every context of utterance. It was then enquired as to what determines the character of 'AS', and it was argued that since 'AS' can be true or false in different languages, and this difference in truth value is not explainable by how the world is itself (since the debate over the truth-value is meant to be a verbal issue), then the only thing that determines the character of 'AS' is how it is used in a particular linguistic community, which is to say that it is a linguistic decision that determines the character of 'AS'.

This has the effect of making of all necessary truths true in virtue of a linguistic decision – so 'LP*' itself and 'AS' are necessarily true in the CSO object language and meta language respectively in virtue of linguistic decisions. What this means is that from the perspective of the CSO-speaker, although there is a *seeming distinction* between analytic truths and synthetic truths in virtue of the fact that some attitude sentences such as 'AS' are true whilst other attitude sentences³¹ will be false, the seeming difference between analytic truths and necessary truths which are synthetic is due only to *another linguistic decision* at a higher level about what attitude sentences are true or false, and so it looks the distinction between analytic truths and synthetic necessary truths is only *nominally maintained* within the CSO-language. This will be equally true for the nihilist also, for they will also hold some object-level necessarily true sentences to be synthetic, such as 'there are no composite objects', whilst maintaining others as analytic. What we find then is that this seeming distinction is merely *illusory* – it is a distinction with no difference, since the necessarily true sentences of all languages are determined as such by linguistic decision. The idea now is that 'analytic' and 'synthetic' are merely arbitrary labels that are used in the meta-languages of each side, and this maintains the illusion of a substantive debate, but in reality the labels denote no genuine difference. One might wonder at this point why this doesn't amount to a rejection of the analytic/synthetic distinction. Against this I would argue that here Hirsch would not be rejecting a distinction, but instead commits to the distinction being *illusory*, and that in reality all necessary truths are analytic. The mistake of the ontologists is to fall for the illusion.

³¹ For example 'The sentence "all bachelors are unmarried males" is synthetic' would be false in the CSO meta-language since it is considered analytic.

I want to now suggest that this treatment of necessary truths appears to commit Hirsch to what I will term Tractarian-analyticity³², or T-analyticity for short. Having analysed what T-analyticity amounts to, I will explain why Hirsch is committed to this thesis, and I will then consider some problems arising from this commitment for Hirsch's position.

Soames characterises T-analyticity in the following way:

...the truth of an analytic statement is supposed to be due entirely to its meaning, whatever facts there may be in the world are irrelevant. *An analytic truth places no constraints on the way the world is, and therefore makes no genuine claim about it, including no claim about what exists in it. Hence, it can't be ontologically worrying.* This is the classical Tractarian doctrine of analyticity, which identifies the necessary and the a priori with the analytic, while maintaining that such statements tell us nothing about the world. (Soames, 2009: 433 [emphasis in the original])

On this view of analyticity then, all a priori necessary truths are analytic, and analyticity is to be understood as being true in virtue of meaning and not telling us anything about the world itself. If T-analytic sentences do not tell us anything about the world, then what do they do? Carnap thought that T-analytic sentences functioned as semantic rules which frame how we are to describe our experiences about the world. Consider that he claims 'to accept the thing world means nothing more than to accept a certain form of language, in other words, *to accept rules for forming statements*, and for testing, accepting or rejecting them' (Carnap, 1950; cited in Soames, 2009: 430 [emphasis added]). The 'thing world' for Carnap, was just a linguistic scheme that we use to describe our experiences as stemming from a world of externally existing material objects. For example, in the 'thing world' language it is a necessary truth that 'the world contains material objects that exist independently of our experiencing them', or something to that effect. Thus any sentence we use to describe an experience must be framed in accordance with that necessary truth, which is to say that no sentence describing an experience can be true if it contradicts a necessary truth, just as no move in a game may be permitted

³² I am not concerned with a historical exegesis of the notion of analyticity as it is presented in the *Tractatus* – it will be seen that I am borrowing the phrase from Soames (2009), who appears to be discussing the concept of analyticity as it was taken over from the *Tractatus* by Carnap, who used the concept to formulate his internal/external distinction.

if it transgresses the rules. This is why necessary truths for Carnap are thought of as semantic rules. Because necessary truths are just semantic rules then, they are true with respect to any possible experience and situation – no possible experience will give us cause to revise what are considered to be a priori necessary truths because necessary truths *are not about the world*; they only determine how it is we are to speak about the world. Thus Soames goes on to add that

We are asked to imagine a choice between our ordinary physical-object framework and a (suitably elaborated) Berkeleyan alternative that speaks only of minds and “sense data”. We are told that this is *simply* a choice between two linguistic schemes for describing experience. There is, we are told, “no belief or assertion or assumption” *in the reality of the thing world* that one adopts when one opts for the thing, rather than the phenomenal framework (Soames, 2009: 430).

Thus according to Carnap, the physical-object linguistic framework and the Berkeleyan framework differ only with respect to the semantic rules that govern how we are to form sentences describing experience within each framework. The semantic rules do not make any claim about the world itself – and in virtue of being semantic rules they come out as being necessarily true by default within the respective linguistic schemes in which they are accepted as such. Therefore according to this view, in accepting the truth of the *thing world*, all we are doing is making the decision to formulate our sentences to describe experience according to a certain set of rules – nothing more. For Carnap, philosophical error occurs when we presume the necessary truths themselves are about the world. To ask whether a necessary truth is true irrespective of a linguistic framework was for Carnap to ask a pseudo question with no genuine answer, and when one adopts a particular linguistic framework, questions about whether a particular sentence is necessarily true or false will be trivially answered by default depending on which scheme we have adopted.

This is the characterisation of T-analyticity that underlies Carnap’s internal/external framework. Let’s say that to accept the premise ‘T-an’ is to accept that

- 1) All a priori necessary truths are analytic
- 2) A sentence is analytic if it is true in virtue of meaning (i.e. is accepted a semantic rule) and does not make any claim about the world itself.

I am now going to argue that Hirsch appears to be committed to T-an. His claim that there is nothing substantively at issue between the nihilist and the CSO can only be maintained if T-an is accepted.

First, remember that for Hirsch the character of a sentence determines its content in a context of utterance. Content is taken to be a set of possible worlds in which the sentence is true. Thus, for Hirsch, the meaning of a sentence can be said to be determined by its character³³. If this is the case, then it looks like for Hirsch necessary truths are true in virtue of meaning. Consider that for a sentence like ‘there is a table in R’, the character of the sentence determines it will have a different truth value in different contexts of utterance – thus whether this sentence is true depends on what context it is uttered in – so it will depend on the world being some way for its truth. Thus it is not true purely in virtue of meaning. However, necessary truths on this view have a character such that they are true regardless of context of utterance. Thus the character of a necessarily true sentence determines that it is true in virtue of meaning. Additionally, of course, character has to be determined by linguistic decision if the debate is to be verbal. If the character of a necessary truth in one language were not determined by linguistic decision, then it looks difficult to maintain that the reason the same sentence is false in another language is merely a verbal issue. Thus for Hirsch it must be that a priori necessary truths are analytic because their character determines that they are true in virtue of meaning, and the reason that the sentence in question has that particular character has to be explainable as arising from nothing but a linguistic decision. Finally, it has to be the case that *all* necessary truths are explained as such. This is because if some necessary truths were not true purely in virtue of meaning, then they would be true in virtue of describing some necessary feature of the world itself, and if that is the case then a disagreement over the truth value of such a sentence would be a substantive disagreement. Therefore it looks like Hirsch’s position commits him to accepting T-an.

6: Does the commitment to T-an pose a problem for Hirsch?

Thus far we have seen that Hirsch’s treatment of a priori necessary truths seems to commit him to accepting T-an. In this section I’m going to discuss a number of issues that arise for Hirsch in light of this. There are two minor issues, and two

³³ Though see chapters five and six for discussion about hyperintensionality.

major issues. The minor issues are that this treatment could potentially lead to an infinite regress of argument between the nihilist and the CSO. However, in one respect, this can be seen as a further indication that the CSO and the nihilist are engaged in a terminal dispute only because they are attempting to maintain that some necessary truths are synthetic. However, this leads to a problem – if T-an is true then must be capable of being asserted as a true sentence of some language. However, it could not be stated as true in any of the meta-languages we have thus far discussed. Insofar as Hirsch maintains that he is a CSO speaker, he cannot be a CSO speaker to the extent that he would have to reject ‘T-an’, stated as a sentence, if he is wearing his CSO hat. But he needs to maintain ‘T-an’ as true in order to explain how the debate is verbal. How Hirsch can reconcile these differences is the second minor issue.

The major issues are that committing to T-an means Hirsch must give up on his claim of being a robust-realist. He must also then be prepared for the possibility that other ontological debates fall victim to T-an, turning out to be merely verbal also. Thus Hirsch must give up on the idea that his merely-verbal claim is only restricted to physical object ontology. Giving up on both these claims means that Hirsch can no longer maintain his modified light version of Carnapianism and must retreat into a thoroughgoing Carnapianism.

The first issue to consider is that the analytic problem may yet re-emerge, albeit at the level of the meta-languages. The original analytic problem was that in order for debate between the nihilist and the CSO to be verbal, ‘LP*’ had to be analytic. Since the CSO-language does not regard ‘LP*’ as being analytic, then ‘LP*’ couldn’t be assumed to be analytic for the purpose of analysing the debate as merely verbal without begging the question against the CSO speaker. The suggested solution in this chapter has been that this problem can be avoided if ‘LP*’ can be analytic from the perspective of the nihilist, but synthetic from the perspective of the CSO-speaker. I suggested that this could be done by accepting that ‘AS2’ could be regarded as true in the CSO meta-language, whilst it could be regarded as false in the nihilist meta-language. Thus from the perspective of each side, ‘AS2’ has a truth value that is in line with their first-order beliefs.

However, it now looks like Bennett can make the same point she made about ‘LP*’ – that it has to be analytic for the debate to be verbal - about ‘AS2’ itself. She can claim that in order for the debate to be verbal, ‘AS2’ has to be analytic (i.e. true in a

virtue of meaning/linguistic decision). However, the CSO-speaker will not regard 'AS2' as being analytic. Remember that 'AS2' is asserting that 'LP*' is synthetic. The CSO will argue that 'AS2' is true because it predicates some property of the sentence 'LP*', and that the sentence 'LP*' does indeed have the property, and that property of the sentence itself will be a result of the semantic properties of the sub-sentential expressions in the sentence and how those expressions relate to objects in the world. Ultimately then the CSO will argue that it is not merely a linguistic decision at the level of the meta-language to label 'LP*' as synthetic. However, as we have no seen Hirsch's solution to the original analytic problem requires this to be the case. Thus it appears that the analytic problem reemerges at the level of the meta-language.

However, at this point an interesting move is available to Hirsch – he can reapply the same reasoning used to solve the initial analytic problem to the problem as it emerges at the next level. Thus he can say that 'AS2' itself only needs to be synthetic (i.e. not true solely in virtue of a linguistic decision) *from the perspective of the CSO*. In order to do this, we can introduce *another* attitude sentence, which is about 'AS2' –

AS3: 'The sentence "AS2" is not analytic'

and we can then argue that 'AS3' can be true in the CSO meta-meta language whilst false in the nihilist meta-meta language. Thus whether 'AS2' is analytic or not is *also* a verbal debate. This quite clearly leads to a regress – as again the analytic problem can be seen to emerge once more regarding 'AS3'. The question is who is this regress bad for? Initially one might argue that if Hirsch's analysis generates this regress then Hirsch hasn't really demonstrated that the debate is merely verbal, since all Hirsch is doing is continually shifting the substantive bump under the carpet without every really getting rid of it. Equally however, it could be argued on behalf of Hirsch that the generation of the infinite regress is a symptom of the initial debate being ill-formulated. As I mentioned earlier, Hirsch can argue that the nihilist and the CSO are falling for the illusion that there is a distinction between necessary truths that are analytic and necessary truths that are synthetic. The reality is that all necessary truths are analytic and if this was accepted, then the regress wouldn't be generated.

One might wonder at this point why Hirsch doesn't simply explicitly make the point

that all necessary truths are analytic. As we have seen in this chapter, Hirsch can avoid the analytic problem if he makes all necessary truths analytic, so why bother with the pretence of maintaining a nominal/illusory distinction in the object-languages between analytic and synthetic necessary truths – why not just deny this distinction outright? I will return to this issue momentarily but I think the ultimate reason is that Hirsch doesn't *want* to have to rely on 'T-an' to establish his merely-verbal claim, but that it very much looks like he has to.

For now let's consider the second issue. Remember now that Hirsch seem to be committed to 'T-an'. But it also seems reasonable to suppose that 'T-an' as a sentence should be true in some meta-language. However, it quite clearly cannot be true in the CSO or nihilist meta-language, since each of those languages will have true sentences that contradict 'T-an'. For example, 'AS2', which is supposed to be a true sentence of the CSO meta-language explicitly states that at least one necessarily true sentence ('LP*') is synthetic. Against this then what can Hirsch do? He might insist that although there are a number of object-languages with attendant meta-languages, there is some privileged meta-language in which 'T-an' is true, and 'T-an' can be true in that language without thereby contradicting what is said in the other meta-languages. One way of understanding this is of thinking of the privileged meta-language as explaining the *reality* of what is going on – i.e. really speaking all necessary truths are analytic, whilst the other meta-languages are in some sense inferior to this, and thus do not really express what is going on. This is quite hard to formulate, and it is not clear how much sense it makes. Against this, as I mentioned, Hirsch might just explicitly state that all necessary truths are analytic and any statement to the contrary, in any language, is false. This will certainly level the playing field since both the nihilist and the CSO will want to maintain that some necessarily true sentences are synthetic, and so insisting that T-an is true will make the nihilist and the CSO equally wrong. However, this is problematic for Hirsch because by committing to T-an, he will have to give up two important commitments. By giving up these commitments, Hirsch will be abandoning his modified light version of Carnapianism and retreating into a thoroughgoing Carnapianism. Let me now discuss what these commitments are.

Unlike Carnap, who envisions that his external/internal distinction will dissolve all metaphysical debate, Hirsch envisions his merely-verbal claim as applying only to debates in physical object ontology (Hirsch, 2005: 67). The problem with this is

that if 'T-an' is true, it will apply to *all* a priori necessary truths. Thus any philosophical dispute that turns on a disagreement over whether a particular a priori necessarily true sentence is analytic or synthetic will turn out to be verbal. Consider, for example, someone who thinks that there are only brain states, but that all statements about mental states are disguised statements about brain states. On this view the sentence about the brain state ('B') would be intensionally equivalent to the statement about the mental state ('M'). Consider someone else who holds that brain states and mental states are distinct, but that there is necessary connection between particular brain states and mental states. Again on this view, a particular statement about a brain state ('B') and a particular statement about a mental state ('M') would be intensionally equivalent. In the former case then one can assert that 'necessarily if B then M', but one can equally assert this in latter case as well. The only difference is that in the former case the sentence would be regarded as analytic, but in the latter case it would be regarded as synthetic and as pointing out a necessary connection in the world itself (between mental states and brain states). However, given T-analyticity this sentence could *only* be analytic. Thus the seemingly substantive dispute about the ontology of mental states would end up being merely a verbal dispute on this approach. Whilst Carnap would welcome such a result, Hirsch does not envision his merely-verbal claim as extending to such debates.

The second claim that Hirsch cannot maintain I think is more serious. Remember from chapter two, section 7 that Hirsch opposes himself to Carnap as a 'robust' realist, whereas he construes Carnap as an anti-realist and/or a verificationist³⁴. Just what it is that makes someone a realist can vary, but here Hirsch takes a central tenet of realism to be an acceptance of the following premises:

- a) The world and the things in it exist for the most part in complete independence of our knowledge or language
- b) Our linguistic choices do not determine what exists

I think Hirsch can coherently maintain premise *b* to the extent that once the character of an existential claim has been determined in a language, then it is true *in that language* to say that linguistic decisions do not determine what exists, for whether an existential claim in that language is true or false is not determined

³⁴ I have already shown that one purported solution to the analytic problem (chapter three, section 6.1) pushes Hirsch toward verificationism, and toward a thorough going Carnapianism.

solely by linguistic decision but will also depend on what contexts the existential claim is uttered in. For example, the sentence ‘there is a table’ has a different character in the CSO-language and the N-language, and this I maintain is a result of a linguistic decision about how that sentence is used in the respective languages, but it is not then purely a linguistic decision that determines that ‘there is a table’ is true or false in the CSO-language, for it will depend on what context the sentence is uttered in, and thus its truth or falsity in a context of utterance will depend on some condition in the world being fulfilled or unfulfilled.

However let us now focus our attention on premise *a*. Premise *a*, as a sentence, has to be stated in some language. And it seems to me that a *robust* realist will want to regard *a*, as a sentence, as *not being true purely in virtue of a linguistic decision*. For if it was simply true in virtue of a linguistic decision, then it could equally be necessarily false in virtue of a linguistic decision. In that case it looks incoherent to assert that the world and the things in it exist for the most part in complete independence of our language if that claim *itself* only comes out true in certain languages. For we could imagine a ‘realist’ language in which *a* is necessarily true (but analytically true), and an ‘anti-realist’ language where *a* is necessarily false (analytically so³⁵). However, to concede that as a possibility seems to concede that one is a realist in a robust sense – at the meta-ontological level one would have to admit that the realism/anti-realism debate itself is merely verbal. The only way to deny that the anti-realism/realism debate (construed as a dispute over the truth of *a*) is *not* verbal is to reject ‘T-an’ from the meta-ontological perspective. However, if Hirsch wants to maintain that debates in physical object ontology are merely verbal, then he cannot reject ‘T-an’ from the meta-ontological perspective. In which case the *only* sense in which Hirsch can be a realist is that he ordinarily speaks the CSO-language, in which *a* is regarded as necessarily true.

This cannot be seriously construed as a robustly realist stance without obstinate insistence, because what determines that *a* is true is a linguistic decision (i.e. the decision to speak a ‘realist’ language where *a* comes out necessarily true). Thus the truth of the claim that the world contains objects independently of language *depends on the language we speak*, which looks incoherent. Against this of course, Hirsch could make the move suggested earlier in this chapter – he could argue that the CSO-language attitude sentence which states that *a* is not analytic comes out

³⁵ In the sense that the negation of *a*, in the anti-realist language, is a necessary truth – i.e. a semantic rule of the anti-realist language.

true in the CSO-language. However, this itself would appear to generate the same regress generated earlier. It seems to me that a genuine robust realist should not concede the possibility of this regress. A robust realist should, at the meta-ontological level, deny ‘T-an’, for that is the only way one can deny that the realism/anti-realism debate itself is verbal. Since Hirsch cannot do that, I argue that he cannot maintain *a* in any serious sense such that it distinguishes his position from Carnap’s. Consider that Carnap has remarked that ‘the Circle³⁶ rejected both the thesis of the reality of the external world and the thesis of its irreality as pseudo-statements’ (Carnap 1950, cited in Price 2009: 324). In other words, for Carnap, the question of whether the world is real or not *itself* is an external question, it has no answer independently of the linguistic framework we chose to adopt.

I think that to be a robust realist one must reject this outlook, and thus the only kind of realism open to Hirsch is a kind of *watered down* realism (whereby he argues that the language he is speaking is one where *a* is true). If Hirsch is forced to concede that the thesis of the mind/language independent reality of the world itself is merely a verbal debate, then by proxy he is also forced to give up any pretence that he is defending common sense in any robust sense, since common-sense takes it that the world does indeed exist in independence of thought and language. Thus to the extent that Hirsch thinks he is defending common sense, he is only doing so in a superficial sense, since no robust defence of common-sense could allow the possibility that the notion of the world’s independent existence was merely a matter of which language we elected to speak.

Thus if Hirsch takes his solution to the analytic problem in the direction of this chapter, he must ultimately commit to the claim that all a priori necessary truths are analytic, and that then means that Hirsch must give up two key commitments which he intended to use to distance himself from a thoroughgoing Carnapianism. As I have shown here then, Hirsch *cannot* solve the analytic problem without retreating into a thoroughgoing Carnapianism, further demonstrating that his position, as it is presented is, untenable.

³⁶ Referring to the Vienna Circle.

7: Summary of chapter four

In this chapter I explored two potential solutions to the analytic problem. The first was to argue that Hirsch need not take a stand on whether ‘LP*’ was analytic or not – he just needed it to be possible for one side to plausibly interpret it as such. Against this, reiterating the argument from the previous chapter, I argued that Hirsch cannot ignore what ontologists have to say about meta-level semantic propositions since ontological debate will manifest itself at this level also – i.e. ontologists will have a stake on what meta-level semantic propositions are true.

The next solution then was to argue that ‘AS’ (i.e. ‘the sentence ‘LP*’ is analytic’) could have different truth values in different meta-languages, and that each side could be understood as speaking the truth in their own meta-language as well as their own object-language. This ultimately would make the dispute over ‘AS’ itself another verbal debate.

I then argued in section four that this ultimately means that all necessary truths turn out to be true solely in virtue of linguistic decision to Hirsch. This, I argue, committed Hirsch to premise that all necessary truths are analytic, in the sense that they are true solely in virtue of meaning and do not tell us anything about the world itself.

This commitment was then shown to be problematic for Hirsch for a number of reasons. The major problem is that it forces Hirsch to give up two important commitments – the claim that he is a robust realist, and the claim that his merely-verbal claim only applies to debates in physical object ontology.

Let’s now take stock of the problems presented to Hirsch thus far. The analytic problem has brought into sharp focus the importance of the analytic/synthetic distinction for Hirsch. It is quite clear from the discussion in chapter three that Hirsch cannot allow a clear-cut, substantive difference between necessarily true sentences that are analytic and those that are synthetic. Thus he must, if he is to avoid the analytic problem, take some kind of stand on the analytic/synthetic distinction in a way that allows him to maintain his central claim that the debates in physical object ontology are merely verbal. However, what I have shown thus far is that the options open to Hirsch that would allow him to maintain his central claim force him to give up a number of commitments. These commitments, to my

mind, are important to Hirsch's vision as he wants his position to be Carnap-inspired, but he certainly doesn't want his position to be thoroughgoing Carnapianism. The fact that Hirsch cannot avoid such a retreat in thoroughgoing Carnapianism if he is to maintain his central claim demonstrates that his position, as he envisions it, is untenable.

Chapter five: Hawthorne's criticisms of Hirsch

1: Introduction

In this chapter I now move in to discuss several criticisms made against Hirsch by Hawthorne. Hawthorne's criticisms stem from what he calls Hirsch's 'intension-centric' approach to semantics (Hawthorne 2009: 225). I shall explain what Hawthorne means by an 'intension-centric' approach to semantics and discuss the particular problems that Hawthorne thinks this approach creates for Hirsch's semanticist claims. These are criticisms are divided into two 'intensional worries', and a more general discussion about hyperintensionality³⁷ and how this impacts Hirsch's semanticism.

I will show that Hawthorne's two intensional worries are not effective criticisms of Hirsch's position, and will then argue that whilst Hawthorne is right to point out that hyperintensionality is potentially problematic for Hirsch, he does not do nearly enough in demonstrating exactly *why* it is a problem for Hirsch's position. This will pave the way for chapter six where I develop a criticism of Hirsch's position, showing a further way in which Hirsch's explanation of how the debates in physical object ontology ignores the ways in which ontological positions have a stake in meta-level claims.

2: The intension-centric approach to semantics

As is by now familiar, for Hirsch a language is a distribution of sentential characters over syntactic structures. The character of a sentence is a function from contexts to intensions, and the intension of a sentence is a function from possible worlds to truth values. The content of a sentence is then understood to be the set of possible worlds in which the sentence is true. Two sentences which have the same character will be intensionally equivalent in every context of utterance, and thus will express the same content in every context of utterance³⁸. On this approach, the sentences 'there is a table' and 'there are simples arranged table-wise' will express the same content in every context of utterance because they are said to have the same character. It is because of this that Hirsch insists that there is no difference between speaking the N-language and the CSO-language. He

³⁷ Briefly, a hyperintensional operator is such that one can replace a true sentence within the scope of a hyperintensional operator with an intensionally equivalent sentence and generate a shift in truth value. Thus intensionally equivalent sentences can be said to differ in their hyperintensional content. See chapter six for more on this.

³⁸ See chapter two, section 4.

argues that describing the world using the N-language is just as good as describing the world using the CSO-language because neither side can assert a true sentence about the world for which the other side cannot supply an intensionally equivalent sentence.

As such then, Hawthorne describes Hirsch's position as follows:

Hirsch...tend[s] to operate with a picture according to which standard ontological frameworks are all legitimate means of describing the same set of possibilities and hence are all intensionally adequate (Hawthorne, 2009: 220).

Thus he suggests that for Hirsch,

different ontologies are to be regarded as countenancing the same set of possible worlds, albeit described using different linguistic schemes (Hawthorne, 2009: 220).

Thus for Hirsch, different ontological positions are merely a *relabeling of the same possibilities*, and for that reason there is no genuine disagreement going on about what the world is actually like independent of language.

3: The first intensional worry

Here Hawthorne offers his first criticism of Hirsch's approach. First he introduces the notion of *intensional advance*. One side in a debate intensionally advances over another if the former accepts true intensions that the latter does not accept (Hawthorne, 2009: 220). Although Hawthorne doesn't really spell out this notion out, what he seems to have in mind is the following idea: The CSO will be happy to accept that they do not intensionally advance over the nihilist in the sense that for any true sentence they can produce, the nihilist can produce a true sentence of their own language which is intensionally equivalent to it – the idea then is the N-language is at least as expressively sufficient for describing the world as the CSO-language³⁹. This then forms the basis of Hirsch's claim that the CSO and the nihilist do not actually disagree with each other about what the world is like. However, Hawthorne argues that this ignores a crucial point. Although the CSO will be happy to accept that he expresses no true intensions that that the nihilist does

³⁹ But only the level of intensionally equivalent sentences. In the next chapter I consider the idea of there being structural features of reality that can only be expressed using hyperintensional operators.

not also express, he can regard his own theory as superior because he will avoid expressing *false* intensions⁴⁰, such as the necessarily false⁴¹ ‘there are only simples’. As Hawthorne says ‘one may express no extra true intensions but still avoid expressing certain false ones’ (Hawthorne, 2009: 221). As such, Hawthorne claims that ‘the claim of no intensional advance as between standard ontological views is thus a rather limited one vis-à-vis the overall project of superficialism⁴²’ (Hawthorne, 2009: 221).

Hawthorne’s first point then is that although one side may not intensionally advance over another, the former may still have a superior theory because they do not accept any false intensions either, whereas the latter accepts false intensions. In the sense that it is better to accept true intensions and reject false ones, the former is doing better at enquiring into the world. By ignoring this, Hirsch seems to ignore an important way in which theories can differ even if neither side intensionally advances over the other.

There is an obvious problem with this criticism however – it ignores the whole idea of charitably interpreting the other side’s sentences so they come out true in that language. This is a *fundamental* part of Hirsch’s position. Let’s take Hawthorne’s problematic sentence ‘there are only simples’. This, he suggests is false, but the nihilist regards it as true and in that sense his theory is inferior to the CSO’s. However, in chapter two section 6.2 it was shown that the CSO can charitably regard the nihilist quantifier as being semantically restricted to range over only simples. In that case, the sentence ‘there are only simples’ comes out true in the nihilist language and thus, *once charitably interpreted*, the nihilist cannot be said to be accept a false intension, and thus cannot be said to intensionally regress from the CSO. To not pay heed to this central aspect of Hirsch’s position renders this criticism ineffective and shouldn’t really cause any concern for Hirsch. Thus the

⁴⁰ It is not entirely clear to me what Hawthorne means by a false intension. Intensions are not the *bearers* of truth-values; it is sentences that have the property of being true or false on this approach. We might say then to accept a false intension is to falsely suppose that the intension of S delivers TRUE at *w* wherein fact the intension of S actually delivers FALSE at *w*.

⁴¹ Hawthorne considers this necessarily false, though it may not be necessarily false if there is a world containing a number of highly scattered simples – the CSO can countenance a world of scattered simples that do not compose anything because according to the CSO, composition requires that the simples be arranged or joined in a certain way. As such a better example would be ‘there are no composite objects’.

⁴² Hawthorne describes Hirsch’s meta-ontology as ‘superficialism’.

first intensional worry that Hawthorne raises shouldn't be seen as a problem Hirsch.

4: The second intensional worry

However, Hawthorne presents a more problematic 'intensional worry' for Hirsch. He argues that the no-intensional advance claim is misguided because 'it forgets that certain ontologies multiply possibilities in ways that are resisted by other ontologies' (Hawthorne, 2009: 221). This is problematic for Hirsch for the following reason: if he insists that all ontological positions simply re-label the *same set of possible worlds* this will infringe on the first order debate in a non-neutral way. This is because one side might recognise n distinct possibilities (and thus n possible worlds), whereas another side might recognise $n+1$ distinct possibilities – as such different ontologies will 'shape possibility space in different ways' (Hawthorne, 2009: 230). If theory T recognises $n+1$ distinct possibilities, and theory T^* recognises n possibilities, then proponents of T will think they intensionally advance over proponents of T^* , and that will be part of the *first order theory*. The very fact that they disagree over the shape of possibility space thus forms part of the disagreement. Thus if Hirsch insists that T and T^* simply re-label the same possibilities, then either proponents of T will be *wrong* in thinking that they intensionally advance over T^* , or proponents of T^* will speak in a way that is too intensionally coarse-grained to distinguish genuinely distinct possibilities. Either way, the debate is no longer then verbal, because for the debate to be verbal each side ought to be able to recognise that they do not intensionally advance over the other. Thus if it is part of the first order disagreement that one side *does* think they intensionally advance over the other, then Hirsch is going to have trouble establishing his claim that the disagreement is merely verbal in his intended sense, as Hawthorne points out:

...the claim that a view intensionally advances over another depend[s] on whether it is true: one thus cannot evaluate the claim that it distinguishes possibilities that the other is blind to from a neutral perspective (Hawthorne 2009: 224).

Thus imagine that Hirsch says in fact T and T^* are intensionally equivalent and T does not therefore intensionally advance over T^* , this means that he is actually taking sides in the first order debate on the side of proponents of T^* , who thought all along that proponents of T were incorrectly positing distinct possibilities where

none existed. On the other hand, if Hirsch says for example, that *T* *does* indeed intensionally advance over *T**, then the debate is no longer verbal then either, because theories will no longer be intensionally equivalent. Thus it looks like an ontological disagreement that results in a disagreement over the shape of possibility space is one that can't be verbal in the sense intended by Hirsch. Hawthorne thus proceeds to discuss four examples of ontological disagreements where this indeed appears to be the case:

- i) The statue vs. the lump ontologist
- ii) The nihilist vs. the anti-nihilist (the CSO)
- iii) Gunkism vs. atomism
- iv) Endurantism vs. perdurantism

I think the general problem that Hawthorne has presented here is a good one, but the problem I have is with the examples he has used – I do not think the examples he focuses on are problematic for Hirsch, especially in the manner that Hawthorne has constructed them, because Hirsch's meta-ontological criticism is *narrowly focused* on a handful of first-order disagreements and he has made it well known that he is not a global superficialist/deflationist about ontological debates in general⁴³ (Hirsch, 2005: 67). As such I think the narrow handful of theories that Hirsch does focus on are not really amenable to this objection *because* they are theories where proponents do not claim to intensionally advance over the other and do not disagree over the shape of possibility space. Let's then discuss the four examples and highlight in each case why the particular example is not problematic for Hirsch.

4.1: The statue vs. the lump ontologist and the nihilist vs. the anti-nihilist

Hawthorne imagines two ontological theories – one of particles and lumps, and one of particles lumps and statues. Call the second ontologist the statue-ontologist and the first the lump-ontologist. The statue-ontologist argues that the statue and the lump are numerically distinct objects, whereas presumably the lump-ontologist will argue that there is merely a statue-shaped lump. Hawthorne argues that the statue-ontologist can imagine two worlds which differ *de re* at the level of statues but which are the same *de re* and qualitatively indiscernible at the level of lumps and particles:

⁴³ Though if he accepts the proposed solution to analytic problem from the previous chapter it looks like he might have to give up this distinctive aspect of his position.

It is not hard to imagine the second philosopher allowing for distinct possible worlds that have the same profile when it comes to particles and lumps but different profiles when it comes to statues. After all, he will notice that the same lump can be used to make two different statues – first one, then another. He can then, for example, imagine one world where Ben is the last statue made, Bill the penultimate and infinitely many before them, and another where Bill is the last statue to be made but where there is duplication of the first world when it comes to particles and lumps (Hawthorne, 2009: 221).

He goes on to add that

If one distinguishes the Bill and Ben possibilities one will be confident that, in the sense described earlier, one has intensionally advanced over the particle and lump ontologist. On no vaguely natural translation scheme will one think that that the lump and particle ontologist can express the intension expressed by ‘there is a world which ends with Bill and not with Ben’ (Hawthorne, 2009: 221).

There are two important points to note about this example. First of all the idea that there is an *infinite* series of statues stretching back in time is important to this example, for if the series were only finite, the lump-ontologist could identify the distinct possible worlds by *the number of times the lump had been statue shaped*. If the series is infinite there is no sense that can be made of the idea that in one world there was one additional time that the lump was statue shaped. Secondly, it is important that each statue is qualitatively identical to every other statue down to the position of every single atom. Thus there are no properties that can serve to distinguish Bill from Ben save for their *haecceity* – a property which is not grounded in any qualitatively discernible differences between Bill and Ben⁴⁴. As Hawthorne says, the statue ontologist’s theory will make for a wider range of ‘haecceistic possibilities’ (Hawthorne, 2009: 221).

Here then Hawthorne has given an example of a dispute where one side, the statue-ontologist, will recognise distinct possibilities due to their ontological beliefs.

⁴⁴ The notion of a haecceity, as a property, is such that *a* in *w1* and *b* in *w2* can be identical in every respect – causal origin, history, qualitative properties at any point in time, modal properties, and yet in virtue of the fact that they have different haecceities, *a* ≠ *b*. Thus there can be two worlds which are identical in every respect except for the fact that in one world *a* exists and in the other world *b* exists.

Because of a conflicting ontological belief, the lump-ontologist will not recognise the distinct possibilities. Thus, according to Hawthorne, the statue-ontologist will think they accept more true intensions than the lump-ontologist and thus will argue that they intensionally advance over them. Conversely, the lump-ontologist will regard the statue-ontologist as accepting certain false intensions about what distinct possibilities there are – they will think the statue-ontologists posits distinct possibilities that do not exist.

If Hirsch then argues that the statue-language and the lump-language are intensionally equivalent and thus merely re-label the same set of possibilities, he will infringe on the first order debate. This is because only if the Bill and Ben worlds are *not actually distinct possibilities* can we say that the statue-language and the lump-language are intensionally equivalent. This is because if there are indeed distinct possibilities, then the lump-language will not be intensionally fine-grained enough to distinguish those possibilities (thus the languages will not be intensionally equivalent and the statue-ontologist will be right in thinking he intensionally advances over the lump-ontologist). Thus it cannot be the case that the Bill and Ben worlds are distinct. But in that case the statue ontologist loses the debate! As such, the claim that the statue-language and the lump-language are intensionally equivalent can't really be sustained, and hence the idea that any disagreement between them is merely verbal in Hirsch's intended sense cannot be sustained either.

Hawthorne argues that the same lessons from the statue/lump ontologist example can be applied straight-forwardly to the debate between the nihilist and the anti-nihilist (we shall retain the nihilist and the CSO terminology) – for he claims that 'If one accepts the possibility of pairs of worlds that differ de re at the level of certain macro objects but are alike – qualitatively and de re – at the level of simples, then one will think the nihilist's language intensionally too coarse grained' (Hawthorne, 2009: 222). If he is right about this, then it does look like a Hawthorne has found a serious problem with Hirsch's position because the dispute between the nihilist and the CSO is one of the disputes that he regards as being merely verbal.

However, the way that Hawthorne has constructed the debate between the statue and the lump ontologist, and the way he has construed the debate between the nihilist and the CSO are problematic.

Let us regard the particular way that Hawthorne has constructed the lump and statue ontologist debate. I think the general issue of whether there are two distinct entities – a lump and a statue, or just a statue-shaped lump, is one which Hirsch himself would find a trivial verbal dispute. However, the way Hawthorne has set up the example is that Bill and Ben are all qualitative duplicates such that there is no discernible difference we can use to ground their unique identity. The other option would therefore be to appeal to historical or relational properties. For example, we might identify Bill as the n th statue in the series whilst Ben is the n th+1 statue in the series. However, by making the series of statues infinite Hawthorne has also removed this way of grounding the unique identities of the statues across different worlds. As such, Hawthorne suggests that we could ground the different identities of Bill and Ben in virtue of their different *haecceities*.

I think the problem with this is that one can be a statue-ontologist without thereby committing to haecceitism. The thesis that there are qualitatively indiscernible but distinct possible worlds is a separate first order theory that statue-ontologists will not necessarily subscribe to. For example, it could be argued by statue-ontologists that the principles by which we individuate statues should include their historical properties, such as when and by whom they were created. Thus the uniqueness of a statue does not need to be grounded in any intrinsic properties of the statue itself – certain extrinsic properties could be essential to its identity⁴⁵. As such a statue-ontologist that denied haecceitism would proclaim that there is no distinction between the Bill and Ben worlds because there is no way of grounding the identities of the last statues that exist at each world as being non-identical. Consider that Bill, in Hawthorne's example, is taken to exist in both worlds – in the Bill world he is the last statue but in the Ben world he is the penultimate statue. However, a statue ontologist that denied haecceitism would argue that there is no way of identifying these statues as both being Bill. This is because all the statues are qualitatively identical, and thus the only recourse would be to argue that each statue is identical with the statue to which it has a one-to-one correspondence with in the infinite series of statues – in which case the last statues in each world are identical. Thus, anti-haecceity statue ontologists would not recognise this example

⁴⁵ For example, if we have two different worlds containing qualitatively identical objects, where one was formed by the conscious decision of an artist, whereas in the other world the object was formed after a rock was struck by lightning, the statue ontologist might argue that although qualitatively identical, these objects are not identical due to their different causal origins. Indeed it might be argued that as an artefact, part of the identity conditions of a statue include that it was fashioned by a conscious agent.

as being a case of genuinely distinct possibilities – they would argue that possibility space does not contain these kinds of qualitatively identical but distinct worlds. Thus I would argue that one can be a statue ontologist without accepting haecceitism – and *this* is the kind of ontologist that Hirsch would level his superficialist claims against. It has been remarked⁴⁶ that there are some debates that Hirsch thinks are substantive, and he might very well consider the dispute over haecceities such a debate. In my mind however, I think Hirsch would actually regard the thesis that there can be qualitatively indiscernible yet distinct possible worlds as an obviously false first order position, given his general inclination toward common sense. In that case, Hawthorne's example wouldn't be troubling to Hirsch at all.

In summary then, the first example isn't troubling for Hirsch because Hawthorne has constructed it in a way that requires haecceities, which do not fall within Hirsch's purview when considering what sorts of ontological debates are merely verbal.

This leads me on to the next point – Hawthorne argues that the main points made in the lump/statue example can be straightforwardly applied to the nihilist/CSO debate (see above). However, it does not seem to me that this is the case, because once again CSO's *do not normally require haecceities for their position* – I do not think a regular CSO would countenance qualitatively and *de re* identical worlds at the level of particles which differed *de re* at the level of composite objects. This is because generally, the identity of a composite object will be grounded, to a large extent, in its composition. Thus a CSO will argue that worlds can differ qualitatively and *de re* at the level of particles whilst being *de re* identical at the level of macro-physical objects – this is what grounds the persistence conditions of macro-physical objects which can survive a loss or change of parts. For example, let's say a statue called Hugo exists in w^1 and w^2 , and in w^2 Hugo loses a single atom from the top of his head – let's imagine this lonesome atom floats away into space and doesn't fuse to create any new composite objects. According to the CSO, these worlds differ qualitatively and *de re* at the level of atoms but not *de re* at the level of what macro-physical objects exist. They will say Hugo exists in both worlds, but that in one world he has lost an atom. In cases such as this, the nihilist can obviously also recognise these distinct possibilities because the simples arranged

⁴⁶ See this chapter, and chapter four, section 6.

Hugo-wise in $w1$ will actually be a different set of simples to the simples arranged Hugo-wise in $w2$.

It would be a very strange CSO indeed who argued, for example, that worlds which were identical *de re* at and qualitatively at the level of simples differed *de re* concerning what composite objects exists. The notion of haecceities is once again an additional first order belief that a CSO will not necessarily (and in my opinion is unlikely to) subscribe to. As such, Hirsch can concede that the debate between a pro-haecceity CSO and the nihilist might very well be substantive, but that is not the debate he is focused on, and since he is not a global semantic deflationist he can, once again, simply deny the thesis of haecceitism, in which case again Hawthorne's second example would not be troubling for Hirsch.

To summarise then, the examples Hawthorne has focused on here can be seen to be entirely reliant on haecceities. However, Hirsch doesn't focus his semanticist arguments on disputes between those who accept haecceities and those who reject them. If we remove the notion of haecceities, the statue/lump debate and more importantly the nihilist/CSO debate can be shown to be merely verbal in Hirsch's sense. Only with the injection of haecceities can Hawthorne prove his point, so whilst Hawthorne has shown that *some* debates can't be verbal, this shouldn't prove to be decisive against Hirsch, who focuses his verbal claim on a very specific set of debates.

4.2: Gunkism vs. atomism

Proponents of gunky matter argue that matter has no ultimate parts which are themselves not composed out of smaller parts – there is no 'bottom layer'. Opposed to this will be the atomist who will maintain that at the bottom level, matter consists of simples which are not composed of anything smaller – as such they think there is a bottom layer of fundamental 'building blocks'. Nihilists necessarily are atomists in this sense – one can also be a CSO and an atomist in the sense that although one accepts the existence of composite objects, reality ultimately has a bottom layer of simples. It is also the case that someone might be a CSO and might believe in gunky matter. Concerning the debate between the gunkist and the atomist then, Hawthorne argues that:

Gunky matter is directly relevant to the issue as to whether anti-nihilists make intensional advances over mereological nihilists. If one

thought the world was necessarily gunky then unless one went in for absurdly conciliatory semantics, one wouldn't at all think that the nihilist had the resources for expressing the intension expressed by 'There are tables.'⁴⁷ Meanwhile, if one thought that gunk was a contingent possibility then it would similarly be very difficult to find a sentence in the nihilist's mouth that expresses the intension expressed by 'There exist tables,' since one would think that at some worlds tables are gunky (Hawthorne, 2009: 223).

The problem with this example is that Hirsch himself acknowledges that this debate is not verbal⁴⁸ – and so again it is not one that he considers as falling with his purview. It is worth mentioning however, that if matter is gunky then the nihilist cannot have his simples – for whatever is identified as a simple (i.e. an object without proper parts) will always be found to be further composed of proper parts. Thus the debate between the nihilist and the CSO is only verbal if gunkism is necessarily false. In that case then, if Hirsch does indeed insist that the debate between the nihilist and the CSO is merely verbal, he must take the first-order position that gunkism is necessarily false. This is a potentially problematic issue for Hirsch but one which I will not expand up on here. Suffice to say in the context of what Hawthorne is trying to prove, the gunk vs. atomism debate is a poor example of showing Hirsch up seeing as he does not regard this issue as merely verbal.

4.3: *Endurantism vs. Perdurantism*

The final example that Hawthorne gives is the classic thorn in the side of the perdurantist: Imagine two possible worlds, each empty apart from the fact that each contains a lonely, homogenous disc. In one world the disc is spinning about its central point (spin world) and in the other world the disc is stationary (no spin world). Hawthorne argues that

It is easy enough for many kinds of endurantist to distinguish a homogenous stationary disk from one that is spinning – this will turn on facts about whether the constituent particles take a helical

⁴⁷ Because the normal way out for a nihilist here would be to say that 'there are simples arranged table-wise'. However, if there ultimately are no simples however, then it doesn't look like this will work.

⁴⁸ See Hirsch 2005: notes 2 and 29, and Hawthorne 2009: note 15

trajectory through spacetime, where such facts are grounded by the identity of particles over time. But the standard perdurantist countenances so many objects that when confronted with any persisting disk he will recognize particle-sized objects with a helical trajectory and particle-sized objects with a trajectory of a stationary disc. Hence he cannot look to the trajectory of particle-sized objects to ground a distinction between a spinning and stationary disk...this will make for depleted possibility space from the perspective of the perdurantist (Hawthorne, 2009: 223).

The classic problem here then is that this intuitive difference cannot actually be captured by the perdurantist. His ontology makes it impossible to ground the distinction between the possibilities, thus it looks like the endurantist will recognise a wider range of possibilities than the perdurantist. This again is potentially worrying for Hirsch because the endurantism/perdurantism debate is one that he considers to be merely verbal.

The main problem with this example is that it should not be so readily assumed that the lonely homogenous spinning presents an insurmountable problem for the perdurantist position, as there are arguments which suggest that the perdurantist *can* recognise the distinct possibilities. For example, Butterfield, arguing from general relativity, says that:

...the trajectory of a test-particle falling towards a massive body depends on whether (and how) the body is rotating: the rotating mass 'drags', albeit very slightly, the inertial frames in its vicinity (Misner et al. [1973], pp. 699, 879, 1117). This frame-dragging means that the [rotating disc argument] fails in the sense that, in the usual version, the inertial frames (the worldlines of test-particles) are dragged around the rotating disc...but not around [the stationary disc]... In short, the [rotating disc argument] fails because there is a dynamical effect of rotation, to which a 'sufficiently naturalist' perdurantist can appeal so as to answer the challenge of distinguishing the possibilities (Butterfield 2006: 18).

Whether or not one accepts Butterfield's line of reasoning, the possibility remains that there are moves available to the perdurantist that allow them to distinguish between the spin-world and the no-spin world. In general, the kinds of debates that

Hirsch focuses on do not tend to multiply possibilities in different ways. As we can see in this example, the spinning disc thought experiment presents a *problem* for the perdurantist that he would like to get rid of. Thus the perdurantist *aims* at intensional equivalence with the endurantist – he doesn't disagree about the shape of possibility space in the same way a pro-haecceitist will disagree with an anti-haecceitist. As I said earlier, the problem of the lonely discs is a problem for the perdurantist that he would like to get rid of. Thus when Hirsch argues that the disagreement between the endurantist and perdurantist is merely verbal, he is assuming that there are *no major objections that one side can level against the other at the first-order level in a way that will dispose proponents of one side to change their view*. We might think of it this way – the haecceitist specifically makes claims about the nature of possible worlds –namely that they can differ de re whilst being qualitatively indiscernible, whilst the anti-haecceitist will deny that there are such distinct possibilities. In the endurantist/perdurantist case, the perdurantist is not disagreeing with the endurantist in the same way – the challenge laid down by the endurantist is that the perdurantist must find a way of distinguishing the possibilities or he loses the debate – and we have seen an example here of how the perdurantist might meet the challenge. Thus the perdurantist does not wish to recognise any intensional difference between his theory and the endurantist. As such then, assuming the perdurantist can meet the challenge presented by the lonely homogenous spinning disc, this example will once again not cause an issue for Hirsch.

5: Summary of the intensional worries

Before moving onto discuss some of the issues Hawthorne raises regarding hyperintensionality, let's summarise the two intensional worries that he has presented. The first intensional worry was that just showing that theory *a* does not intensionally advance over theory *b* and vice versa is not enough to show that a debate is verbal, because if it is the case that neither theory intensionally advances over the other, it is still possible that one side accepts false intensions that the other side does not. I argued that, as Hirsch sees it, neither side can be accused of doing this because a fundamental part of his argument is that we should charitably re-interpret sentences that would be *prima facie* false according to our own theory as being true in the language that the other side is speaking. Thus the first worry shouldn't be a problem for Hirsch as it fails to take into account this aspect of his position.

The second intensional worry was that Hirsch's approach to superficialism does not take into account the ways that different ontological positions multiply possibility space in different ways. I have argued that the particular debates that fall within Hirsch's purview for criticism do not multiply possibility space in different ways – this is not the central point of contention between the theories he focuses on. The first two of Hawthorne's examples relied on haecceities, whereas Hirsch's claims are directed at ontological positions which do not tend to subscribe to the notion of haecceities to ground the identities of objects. The third example was simply one that Hirsch does think substantive, and finally the fourth example was one where one side of the debate is trying to insist that there are *no* intensional differences between his position and his opponent's – he is not trying to argue that there are! As such, the final example can't really prove the point that Hawthorne is trying to make.

Overall, the second intensional worry that Hawthorne presents is effective in showing how *some* ontological debates can't be merely verbal in the sense intended by Hirsch, but he is wrong to think that this causes a problem for Hirsch, because Hirsch himself is narrowly focused on showing that only the handful of debates comprising physical object ontology are merely verbal – he is happy to concede that some debates will be substantive.

There is an important final point to be made here. In chapter three we left the *analytic problem* unsolved for Hirsch, and in chapter four I argued that in order to solve this problem, Hirsch would have to accept the notion of T-analyticity. This entailed, however, that certain disputes that are outside of physical object ontology get swept up as being merely verbal, and thus Hirsch's insistence that he is not a global semantic deflationist would look threatened. Given that Hawthorne has presented another serious criticism that can only really be countered by Hirsch insisting that it doesn't apply to the narrow set of debates he is interested in, the T-analytic solution to the analytic problem looks even less attractive now, in light of Hawthorne's argument.

6: Hyperintensional issues

As we have seen, so far Hawthorne's arguments haven't been successful in showing Hirsch's position to be problematic. In his paper Hawthorne moves on to discuss what he calls hyperintensional worries – perhaps then these arguments are more successful in showing Hirsch's position to be problematic. There are two main

points that Hawthorne makes. In the first instance, Hawthorne calls to attention the fact that metaphysicians sometimes use hyperintensional operators in their ‘serious theorizing about the world’ but that on an intension-centric approach to semantics such as Hirsch’s, hyperintensional operators cannot act on the content of a sentence. The main issue with Hawthorne here is that he merely calls attention to this point without specifically developing it into an argument showing why this constitutes a problem that Hirsch needs to deal with. That is a task that I take up in chapter six. The second issue that Hawthorne raises is slightly more developed and indeed does look like it causes a problem for Hirsch. Hawthorne points out a number of cases not related to physical object-ontology where a bottom-up approach to semantics (Hirsch’s intension-centric approach is a top-down approach) is explanatorily illuminating whereas the intension-centric approach leaves us with problems (just what these are will be explained below). Whilst Hawthorne is not precisely clear regarding his own argumentative steps here, it suggests that we should prefer a bottom-up approach to semantics rather than a top-down approach. Hawthorne argues that the success of Hirsch’s project is linked to the success of top-down semantics, which I agree with, and then we can make the argument that if a top-down approach to semantics fails, Hirsch’s project fails, and we have (thanks to Hawthorne’s examples) some reason to think that a top-down approach might fail.

Lets’ consider the first issue then. Hawthorne points out that

Hirsch’s work makes it quite clear that he takes the fundamental unit of cognitive significance to be the intensions of sentences – that is, functions from possible worlds to truth values. The semantic behaviour of a sentence is, from this perspective, adequately captured by its character, which profiles the way that intensional content depends upon context (Hawthorne, 2009: 224).

What Hawthorne appears to be saying here is that on an intension-centric approach what we grasp or understand fundamentally, when we understand a sentence, is its intension – i.e. the function from possible worlds to truth values, which is, as Hawthorne reminds us, dependent upon the character of the sentence. By the phrase ‘cognitive significance’ then, he seems to suggest that what one is in cognitive contact with then is the intension of the sentence. This is quite vague but what it points toward is that in knowing or apprehending one sentence, one does

not gain any knowledge or increased understanding from knowing or apprehending another sentence which is intensionally equivalent to the first. Thus any differences between sentences that do not affect the intension of the sentences are not cognitively significant – they have no bearing on the content of the sentence.

This is all well and good for Hirsch, as Hawthorne points out:

If two theories are characterwise equivalent – that is, if there is a one-one map from sentences of one to sentences of the other that is character preserving – then everything about the world that we could come to know by understanding and accepting one theory in a given context we could just as easily learn by understanding and accepting the other in that context⁴⁹ (Hawthorne, 2009: 225).

Thus far this does not strike me as a criticism but appears to be the entire point that Hirsch is making. Hirsch argues, for example, that the four-dimensionalist will concede that a mereological essentialist's 'description of the physical facts seems...to be as adequate as ours' (Hirsch, 2005: 77). Thus Hirsch seems to be suggesting that there is nothing more we can learn about the world that isn't captured at the level of true intensions – intensionally equivalent theories are equivalent in their informative content. This leads on to what Hawthorne says next – that a superficialist like Hirsch will not accept that hyperintensional operators have any 'central place in our serious theorizing about the world' (Hawthorne, 2009: 225). Why is this a problem?

Metaphysicians are comfortable with the idea that certain necessarily co-existing facts might partake in certain asymmetrical relationships of dependence which cannot be adequately expressed without the use of hyperintensional operators, with the additional belief that there are structural features of reality that can only be expressed by the use of hyperintensional operators. For example, the sentences 'Socrates exists' and '{Socrates} exists' are intensionally equivalent – but metaphysicians will consider '{Socrates} exists because Socrates exists' as true whereas they will consider 'Socrates exists because {Socrates} exists' as false⁵⁰. This

⁴⁹ This seems to confirm the idea that cognitive significance is tied to knowledge of some state of affairs – two sentences do not differ in their cognitive significance if by accepting one sentence, a person's understanding or knowledge is not enhanced by then accepting the other second sentence.

⁵⁰ See Bliss, Ricki; Trogon, Kelly, "Metaphysical Grounding", *The Stanford Encyclopedia of Philosophy* (Fall 2015 Edition).

is because it is thought that the existence of the unit set containing Socrates is dependent for its existence upon the existence of Socrates, whereas the existence of Socrates is not dependent upon the existence of the unit set containing Socrates. Thus swapping intensionally equivalent expressions around certain sentential connectives will generate a change in truth value – suggesting that certain structural features of reality, including metaphysically interesting relations, cannot adequately be expressed without hyperintensional operators.

Hawthorne argues that for the superficialist like Hirsch, hyperintensional features of sentences cannot affect the sentence's content. Content is determined solely by the intension of the sentence, and hyperintensional operators are 'regarded as superficial artifacts of the vehicles by which intensions are delivered' (Hawthorne, 2009: 226). Thus here Hawthorne calls attention to a potentially interesting point of tension between Hirsch and the serious ontologist. I say 'calls attention to' here because that is merely all Hawthorne does. In making this observation about hyperintensional operators, Hawthorne does not really establish why this constitutes a problem for *Hirsch* and not the serious ontologist. For it seems that if Hirsch is right in thinking hyperintensional operators have no place in our serious theorizing about the world, then going solely on what Hawthorne discusses here, this looks *all the more worse for the serious ontologist who does take* hyperintensional operators to play a role in our serious theorizing about the world. What Hawthorne needs to do, and what he fails to do, is to actually turn this around on Hirsch and show why the issue of hyperintensional operators is a problem for *him* and *not* the serious ontologist. Thus as things stand regarding Hawthorne's first hyperintensional worry, we cannot really at this stage say whether this constitutes an issue for Hirsch at all. In the next chapter then I will look deeper into this issue and will provide an argument showing exactly why hyperintensionality poses a serious problem for Hirsch.

Let's now consider the second argument. The argument here rests on two examples whereby a bottom-up approach to semantics is explanatorily illuminating whereas a top-down approach is not, suggesting that a bottom-up approach to semantics should be preferred. If this is the case then it will cause some issues for Hirsch. Let's look at the examples Hawthorne gives then.

Hawthorne points out first that if the character of a sentence is *not* generated by underlying referential mechanisms at the levels of words (which is of course what

Hirsch's position implies – see chapter two, sections 4 – 4.2), then 'we should expect the semantic contributions of singular terms and predicates to be determinate only insofar as they are fixed by the characters of sentences that they make up' (Hawthorne, 2009: 228).

He then argues that this approach loses sight of the distinction between deictic and anaphoric uses of the expression 'he'. He imagines a situation where the sentence 'he is happy' is anaphorically linked to the use of a name in a previous sentence. As it so happens, the person in question is also in front of the speaker. Thus, if contrary to fact, the speaker had used 'he' deictically in that context the intension of 'he is happy' would be the same. Thus in that particular context, the character of 'he is happy' delivers the same intension regardless of whether 'he' was used anaphorically or deictically. Hawthorne points out that if facts about 'he' are only recoverable from the intensions of sentences containing 'he' in different contexts, then in that particular context, the different ways in which 'he' might contribute semantically to the sentence is lost (Hawthorne, 2009: 229). Whilst Hawthorne doesn't mention exactly what lesson we are to learn from this example, it is an example whereby a top-down approach to semantics such as Hirsch's would either ignore or be unable to explain an intuitively plausible difference in mechanism by which a sub-sentential expression can refer in the same context, whereas a bottom-up approach would be able to account for this difference. This suggests then that the bottom-up approach is explanatorily more useful than the top-down approach.

Another problem with the top-down intension-centric approach is that all necessary truths are intensionally equivalent, and hence on this account, all necessary truths have the same content (Hawthorne, 2009: 226). The problem with this is that quite clearly if you believe one necessary truth, *you don't thereby believe them all* – thus with necessary truths it doesn't seem to make sense to say that the intension of a sentence is the fundamental unit of cognitive significance.

Hawthorne suggests that in order to explain the 'cognitive cash value' of a necessary truth then, we must be in a position to know what contingent propositions one can learn from knowing the necessary truth – he claims that 'insofar as any putatively necessary metaphysical truth is to earn its keep it now has to do so by generating contingent insights which, construed intensionally, were not already available' (Hawthorne, 2009: 226). Hawthorne himself is sceptical of such an approach to explaining the cognitive cash value of necessary truths, for he later argues that 'it does not seem very natural at all to think that cognitive

achievements associated with pure mathematics all consist in their capacity to put one in touch with true contingent propositions that one was previously ignorant of (Hawthorne, 2009: 228).

What Hawthorne is pointing out in this example is the fact that if the fundamental unit of cognitive significance is the intension of a sentence, that means we require an unlikely and convoluted explanation of the difference in 'cognitive cash value' of mathematical truths (since all mathematical truths are necessary and hence intensionally equivalent). The appeal of a bottom-up approach then is that we can have a fundamental unit of cognitive significance that is more fine-grained than intensions (this point is slightly obscured by the fact that Hawthorne later says the fundamental unit of cognitive significance is a set of possible worlds (Hawthorne, 2009: 230), not an intension – I will return to this point below). On a bottom up approach we could have a structured intension of a sentence, which can be compositionally analysed in terms of the intensions of substantial components – e.g. the intension of a singular term being a function from worlds to individuals (see Hawthorne 2009, 228 note 24) – this would allow us to explain the cognitive difference between necessarily true sentences, which would include mathematical truths. Thus, the bottom-up approach here seems to be explanatorily useful whereas the top-down approach does not.

What we have here then is two examples whereby a bottom-up approach to semantics is explanatorily illuminating whereas a top-down approach is not. It is not clear what the next argumentative step is for Hawthorne, for the discussion is quite rushed. He argues that the success of Hirsch's program rests on scepticism about the explanatory worth of bottom-up semantics arguing that Hirsch's attitude of 'indifference' toward the debate between the nihilist and the CSO rests on the fact that the nihilist can regard 'there is a table' is true in the mouth of the CSO even if he cannot provide in his own language a compositional semantics for the sentence (Hawthorne, 2009: 229). He then goes on to add that 'were one to concede that bottom up semantics is explanatorily illuminating, then reasonable standards of scientific practice would recommend treating the plenitudinous [CSO] framework as metaphysically revealing, and hence abandoning an attitude of indifference' (Hawthorne, 2009: 230). Let's unpack what appears to be Hawthorne's line of thinking here.

First he is claiming that the success of Hirsch's project rests on the top-down approach to semantics. This is I think is true, for Hirsch wants to be able to say that all that matters for a debate between the nihilist and the CSO to be verbal is that the nihilist can provide, in his own language, a true sentence which is intensionally equivalent to the CSO sentence 'there is a table'. He does not need to give a compositional semantics for that sentence if the fundamental unit of cognitive significance is the intension of a sentence, for in that case there is nothing more revealing or illuminating about the sentence 'there is a table' than 'there are simples arranged table-wise'. This is because the content of a sentence is taken to be a set of possible worlds here (determined by the intension of the sentence). Thus what one is in cognitive contact with, when one understands a sentence, on this approach, is a set of possible worlds where that sentence holds true.

What Hawthorne's examples serve to show is that the bottom-up approach is explanatorily illuminating in cases where the top-down approach is not. If there are two competing theories and one can be explanatorily illuminating in cases where the other is not, then other things being equal, we should prefer the more explanatorily illuminating theory. The result of this then is that we should prefer a bottom-up approach to semantics, and that should extend to matters of ontology as well, if we are to avoid special pleading. What does this then mean for the claim that there is nothing more revealing about the sentence 'there is a table?'

Well that claim now looks false, for a bottom-up approach implies that fundamental unit of cognitive significance is not the intension of the sentence (or the set of possible worlds where the sentence is true). On a bottom up approach, the fundamental unit of cognitive significance would be more fine grained – the intensions of the sub-sentential expressions, for example. On this approach then, if both 'there are simples arranged table-wise' and 'there is a table' are regarded as true and intensionally equivalent, a bottom-up approach would reveal a difference in content/cognitive significance between these sentences, and the N-language would look like it was expressively deficient and unable to express the precise content of 'there is a table'. In this sense, the CSO-language would be, as Hawthorne puts it, 'metaphysically revealing', and we should thus abandon the attitude of indifference.

What are we to make of this argument? Hawthorne offers a way out for Hirsch – suggesting that a more ‘careful defence’ of the top-down approach is needed (Hawthorne 2009: 230). Thus it remains to be seen whether with sufficient ingenuity the explanatory problems that Hawthorne has highlighted against the top-down approach will be overcome. It certainly seems to me that Hirsch’s project is reliant on the success of the top-down approach, and thus it seems that if a bottom-up approach to semantics turns out to be the right approach then Hirsch cannot have his verbal claim. In summary of this point then, Hawthorne has given us some reason for thinking that a bottom-up approach to semantics is better than a top-down approach, and has argued that Hirsch’s project rests on the success of the latter. To that extent we can agree with Hawthorne that if top-down semantics is a failure, then so is Hirsch’s meta-ontology. However, this is not a decisive argument against Hirsch for it is contingent upon the failure of top-down semantics, and Hawthorne has only gone a little way in demonstrating why it might fail.

7: Summary of chapter five

In this chapter I covered the arguments raised by Hawthorne against Hirsch. We saw that these divided into ‘intensional’ problems and ‘hyperintensional’ problems. The first intensional problem was that even if two theories were intensionally equivalent such that one side did not accept more true intensions than the other, this didn’t mean that one side might accept one or more *false* intensions not accepted by the other. This was shown to be an inadequate argument for the simple reason that according to Hirsch we should charitably interpret sentences that would be false in our home language as coming out true in an alternate language – thus we couldn’t then claim that they accepted false intensions. The second intensional problem was that certain ontologies ‘multiply possibility space’ in different ways and thus Hirsch is wrong to construe different ontological theories as merely relabeling the same set of possible worlds. This was shown to be inadequate against Hirsch because the examples Hawthorne used did not fall within Hirsch’s purview (see section 6 for a fuller summary of why). Then I moved on to the hyperintensional worries. I think Hawthorne raises a good point here that the success of Hirsch’s project rests upon scepticism toward bottom-up semantics, and that if a bottom-up approach could be seen to be explanatorily illuminating then it would create a problem for Hirsch’s verbal claim. However, this argument is only effective to the extent that it creates doubt about the adequacy of the top-down

approach to semantics and doesn't conclusively demonstrate that there is no merit to that approach. In that sense, the argument is a worry for Hirsch, but it in no way constitutes a conclusive refutation of his meta-ontology.

That leaves us with the final hyperintensional worry – Hawthorne points out that Hirsch's approach to semantics means that hyperintensional operators cannot be metaphysically revealing, but that metaphysicians suppose that they can be. This is an interesting point of tension that Hawthorne *doesn't do anything with* – he doesn't demonstrate why this is a problem for Hirsch and *not* the serious metaphysician. Thus whilst we can credit Hawthorne with bringing this point to our attention, in the next chapter I am going to fully explore this issue and develop it into an argument against Hirsch, showing just why hyperintensionality poses a problem for Hirsch.

Chapter six: The problem of hyperintensionality for Hirsch

1: Introduction

In section 6 of the previous chapter, we saw that Hawthorne raised an interesting point of tension for Hirsch regarding the notion of hyperintensional operators. According to Hawthorne, for a semantic deflationist like Hirsch, hyperintensional operators can have ‘no central place in our serious theorizing about the world’ (Hawthorne, 2009: 225). This puts him at odds with metaphysicians who think that some metaphysically interesting features of reality, such as the nature of the relations between necessarily co-existing objects, can only be expressed using the notion of hyperintensional operators. However, as I remarked in the previous chapter, Hawthorne does not really develop this into a full criticism of Hirsch. In this chapter then I intend to do that.

As Hawthorne sees it, on Hirsch’s view of semantics the content of a sentence is fully determined by its intension in a context, thus intensionally equivalent sentences in a context express the same content (understood by Hirsch to be a set of possible worlds). However, based on textual evidence I argue that Hirsch is committed to the notion of finer-grained distinctions in content between intensionally equivalent sentences. I will then argue that this difference in content can only be brought out with the use of hyperintensional operators, thus I will suggest that Hirsch must recognise that there can be hyperintensional differences in content between intensionally equivalent sentences. I then argue that there are two explanations for hyperintensionality: either hyperintensionality arises from differences in the world itself that cannot be expressed at the intensional level, or that hyperintensional difference arises because of merely *representational differences* of the *same fact* in thought and language. I will apply this to the debate between the CSO and the nihilist, and will argue that in order for this debate to be verbal, hyperintensional differences in content *have to be merely representational*. I will then argue that this generates what I call the H-problem, which in many respects is similar to the analytic problem discussed in chapter three. The H-problem argues that the CSO will be committed to the truth of a certain meta-level sentence which I will call a H-sentence. However, Hirsch needs that particular H-sentence to be *false* in order for the debate to be verbal. Thus, Hirsch has once more ignored the way in which the truth of an ontological position has a stake in certain meta-level semantic propositions, and not merely in object-level propositions. Thus

Hirsch cannot simply assume the truth-values of these without begging the question against at least one side in the debate. It is question begging because without the relevant H-sentence, the CSO language is actually CSO-alt, but the CSO speaker will explicitly deny that they are speaking CSO-alt. Thus although Hirsch might succeed in showing that the debate between the nihilist and the CSO-alt speaker is merely verbal, he cannot assume that the CSO and CSO-alt languages are identical. The only solution for Hirsch is to argue that the debate about the relevant H-sentence is *itself* merely a verbal dispute⁵¹. Thus we will see that the proposed solution to the H-problem is the same as the proposed solution to the analytic problem. Once again, it will be shown that adopting this solution to the H-problem commits Hirsch to T-an⁵², and all the attendant problems that commitment entails for his position. Ultimately then this chapter demonstrates that there are two routes open to committing Hirsch to accepting T-an, and thus two routes open to show how Hirsch's modified light version of Carnapianism ends up being forced back into a thoroughgoing Carnapianism, showing Hirsch's position to be untenable.

2: Fine grained content in Hirsch

Recall from the previous chapter that Hawthorne argued that for Hirsch

The fundamental unit of cognitive significance [is] the intensions of sentences – that is functions from possible worlds to truth values. The semantic behaviour of a sentence is, from this perspective, adequately captured by its character, which profiles the way that intensional content depends upon context (Hawthorne, 2009: 224).

He also later adds that

[Hirsch] thinks that the content of a sentence is a matter of its intensional contributions at contexts (Hawthorne 2009: 226).

Thus Hawthorne construes Hirsch's notion of content as being fully explained and determined by a sentence's intension in a context. Thus character-wise equivalent sentences will express the same content in every context of utterance. Any other difference between character-wise equivalent sentences then, according to

⁵¹ It is not strictly the *only* solution, but it is the most plausible solution. I will discuss other potential solutions in the relevant section, showing them to be inadequate.

⁵² I.e. a commitment to the claim that all a priori necessary truths are analytic and do not tell us anything about the world. See section five of the previous chapter for more detail.

Hawthorne, would have to be merely superficial and wouldn't affect the content of a sentence. However, Hirsch himself seems to admit of there being fine-grained distinctions in content between intensionally equivalent sentences. Consider the following quotation:

...to the extent that we cannot provide a certain kind of compositional semantics for some language...we will be unable to straightforwardly express in our terms the *fine-grained intentional content* of some of the assertions made in that language...some of their *intentional states* may not be fully expressible in our own language (Hirsch, 2005: 80 [emphasis added]).

Here Hirsch is referring to the problem we covered in chapter two, section 3.2. Consider that the N-speaker can regard his sentence 'there are simples arranged table-wise' as having the same character as the CSO-sentence 'there is a table'. However, he cannot say, in his own language, what the CSO-expression 'table' refers to. Hence Hirsch is saying that when the nihilist asserts 'there are simples arranged table-wise' in context *C*, if the CSO asserts 'there is a table' in *C*, then there is no *factual disagreement* going on – but there is nonetheless a difference in fine-grained *intentional content*⁵³. Thus what Hirsch seems to be suggesting, rather loosely, is that the intentional state of the speaker who asserts 'there is a table' is in some sense distinct from the intentional state of one who asserts 'there are simples arranged table-wise'. Hirsch refers to this as a fine-grained distinction in content. It seems to me that whatever this fine-grained distinction in content amounts to, it must be cognitively significant. It seems to me that if two sentences differ in their fine-grained *intentional content* then this distinction must be cognitively *salient* to some degree, otherwise what sense is there to be made of the idea of a fine-grained distinction in intentional content?

Without going into further detail at this moment then, suffice to say for now that Hirsch appears to recognise two grades of content: two sentences are equivalent in their coarse-grained content if they have the same character, but now we have the additional notion of fine-grained content as well – two sentences which are character-wise equivalent can differ in their fine-grained content⁵⁴. Hawthorne

⁵³ Note 'intentional' with a 't'.

⁵⁴ It would appear at this point then that what matters for Hirsch in establishing that a dispute is verbal is that neither side asserts anything that differs in its coarse-grained content from what the other side asserts. It is unclear at this stage how important fine-

earlier suggested that the fundamental unit of cognitive significance on Hirsch's view of semantics was the intension of a sentence, and that the *semantic behaviour*⁵⁵ of a sentence is captured by its character, which determines the intension of a sentence in a context of utterance. However, that now only accounts for the notion of *coarse-grained content*. It does not adequately capture any difference in the supposed fine-grained content between sentences which have the same coarse-grained content. If there is indeed any sense to be made of the notion of a fine-grained distinction in content between the nihilist sentence 'there are simples arranged table-wise' and the CSO sentence 'there is a table', we might reasonably expect that in some special contexts these sentences differ in their semantic behaviour. If not, then why does Hirsch even mention the notion of there being a difference in fine-grained content between these sentences – if this difference in content does not anywhere result in a difference of semantic behaviour then how is it even a difference in content at all?

One way of prizing apart intensionally equivalent sentences to reveal a finer-grained distinction in content is to see how the sentences behave in the context of hyperintensional operators. An operator is hyperintensional iff substitution of intensionally equivalent expressions within the scope of that operator can generate a shift in truth value (Hawthorne, 2009: 225). The idea then is that the sentences 'there are simples arranged table-wise' and 'there is a table', as intensionally equivalent sentences, might reveal their differences in fine-grained content in the scope of a hyperintensional operator.

There is an initial complication with this suggestion. We are initially judging the difference in fine-grained content between different sentences of different languages. It is part of Hirsch's position that the sentence 'there is a table' has a different character in the nihilist language than it does in the CSO-language. Whilst the CSO-language sentence 'there is a table^{CSO}' is intensionally equivalent to the N-language sentence 'there are simples arranged table-wise^N', the N-language sentence 'there is a table^N' is *not* intensionally equivalent to 'there are simples

grained content is then – but it is to be assumed that Hirsch would regard any difference in fine-grained content as being unimportant in accurately describing the *facts themselves*.

⁵⁵ Hawthorne does not specify what he means but semantic behaviour. Since he suggests that the semantic behaviour of a sentence is captured by its character, I think what he is suggesting is that the character adequately specifies the truth-values of the sentence in different contexts of utterance. Consider that the semantic behaviour of a necessary truth on this account is invariant across different contexts of utterance – this invariance is captured by the character of the sentence, which assigns to it the set of all possible worlds in every context of utterance.

arranged table-wise^N. Indeed the very idea of Hirsch's position is that the same *syntactic structure* can have different characters in different languages. Thus whilst the nihilist will regard 'there is a table^N' as necessarily false owing to the character that sentence has in his own language, he can charitably regard 'there is a table^{CSO}' as being true because it has the same character as 'there are simples arranged table-wise^N'. What we then notice is that in the CSO-language, the *syntactic structure* 'there are simples arranged table-wise' is also present, and in that language it also has the same character as 'there is a table^{CSO}'. Thus the two CSO sentences

'there are simples arranged table-wise^{CSO}'

and

'there is a table^{CSO}'

are both intensionally equivalent in every context to 'there are simples arranged table-wise^N'. Since Hirsch recognises a fine-grained distinction in content between 'there are simples arranged table-wise^N' and 'there is a table^{CSO}', it is reasonable to suppose then that there is also a fine-grained distinction in content *within the CSO-language* between 'there are simples arranged table-wise^{CSO}' and 'there is a table^{CSO}', and we have already discussed how this difference in content might be brought out in hyperintensional contexts. The idea then is that if we discover a difference in hyperintensional content between the sentences 'there are simples arranged table-wise^{CSO}' and 'there is a table^{CSO}' *within the CSO-language*, then this also suggests that the difference in fine-grained content *between* the sentences 'there are simples arranged table-wise^N' and 'there is a table^{CSO}' is also hyperintensional.

It just so happens that there are a number of hyperintensional operators that we can use to see that there is a difference in semantic behaviour between the CSO sentences 'there are simples arranged table-wise^{CSO}' and 'there is a table^{CSO}', suggesting that there is indeed a hyperintensional difference in content between these sentences. Consider first the hyperintensional operator 'it is a fundamental fact that...'

It seems to me that whilst a CSO speaker will accept 'it is a fundamental fact that there are simples arranged table-wise', if we substitute 'there is a table' in front of this operator we will generate a shift in truth value, for whilst a CSO speaker will

accept it is a fundamental fact that there are simples arranged table-wise, they will not accept that it is a fundamental fact that there is a table.

Let's now consider a two-place operator: 'the fact that...depends upon the fact that....'

In the CSO-language whilst

'the fact that there is a table depends upon the fact that there are simples arranged table-wise'

will be considered true, it seems to me that the following will be considered false:

'the fact that there are simples arranged table-wise depends upon the fact that there is a table'

Thus we have two examples here where two sentences that have the same coarse-grained content appear to differ in their semantic behaviour within the scope of hyperintensional operators, suggesting that there is indeed a difference in hyperintensional content between the sentences.

3: *Two accounts of hyperintensionality*

We have thus far determined that there appears to be a difference in hyperintensional content between the sentences 'there are simples arranged table-wise^N' and 'there is a table^{CSO}'. There are two explanations for why intensionally equivalent sentences can differ in their hyperintensional content.

One account of hyperintensionality is that there are certain features of reality that cannot be adequately captured at the level of intension alone, and thus hyperintensional operators are required to properly express these features of reality. As Hawthorne points out, metaphysicians argue that hyperintensional operators 'track structural features of reality' (Hawthorne, 2009: 226). A metaphysician might use two place hyperintensional operators to state the anti-symmetric ontological dependence/priority relations between necessarily co-obtaining facts or necessarily co-existing objects. For instance, since in any world where there exist simples arranged table-wise there will also exist a table⁵⁶, the

⁵⁶ According to the CSO of course.

sentences ‘there are simples arranged table-wise’ and ‘there is a table’ will be intensionally equivalent. However, the fact the sentences are intensionally equivalent doesn’t reveal the deeper structural relations that are thought to hold between the facts/objects. The idea then is that certain distinct but necessarily connected features of reality cannot be prised apart at the level of intensions alone – intensions are not fine-grained enough for us to be able to express this deeper structure, and thus hyperintensional operators are required for us to be able to express these kinds of facts. For short then let’s say that one account of hyperintensionality is that it arises from the structure of the world itself.

An alternative account of hyperintensional content is that it arises purely from the different ways we represent reality to ourselves in thought and language, as Nolan suggests:

One might...see hyperintensionality as due to our representations rather than the world. We could think that the world itself has a structure that could be entirely adequately captured in intensional terms (or perhaps even extensional terms), with hyperintensionality arising in our theories only because of the sorts of representations we use (Nolan, 2014: 155).

According to this account of hyperintensionality, any hyperintensional difference between intensionally equivalent sentences is purely representational – hyperintensional differences do not arise from the structure of the world itself. Rather, that structure is ‘read into’ reality by the kinds of representations we use.

As an example of this kind of hyperintensionality, Jespersen has argued that the statements ‘the glass is half empty’ and ‘the glass is half full’ are intensionally equivalent, and both adequately capture the same fact about the state of the liquid in a beer glass. However, he argues there is a difference in understanding arising from the fact that to fill a glass and to empty a glass are conceptually different – they are *procedurally distinct* ways of arriving at the same state of affairs. Thus he argues that whilst both sentences are equally good at describing the same fact, they are *not* equally good at describing someone’s conceptualisation of that fact, as Jespersen argues that someone may be able to conceptualise a glass as being half full without necessarily conceiving it as being half empty (Jespersen, 2010: 97)⁵⁷.

⁵⁷ It could be argued that the same glass exists in worlds A and B, and that in world A it started out full and is thus half empty, and in world B it started out as empty and is thus

The idea then is that hyperintensionality does not arise from structural features of facts themselves, but arise from the way in which we *conceptualise or represent those facts to ourselves*. It seems reasonable then that the same fact may be represented using intensionally equivalent descriptions where those equivalent descriptions involve distinct *concepts*.

Thus we have seen now that there are two accounts explaining why there appears to be a difference in hyperintensional content between the sentences ‘there are simples arranged table-wise^N’ and ‘there is a table^{CSO}’. The question now is which account does Hirsch require for his position to work, and whether this account causes any problems for him?

4: What kind of hyperintensionality does Hirsch require?

In this section I will argue first that from what Hirsch says he seems to imply that any difference in fine-grained content must be due merely to our different representations of the same fact. I will then show why this *has* to be the case if he is to maintain that the debate between the nihilist and the CSO is merely verbal. However, in the next section I will then show that arguing that the hyperintensional difference in content between the sentences ‘there are simples arranged table-wise^N’ and ‘there is a table^{CSO}’ is merely representational cause further problems for Hirsch.

Consider once more the following quotation from Hirsch:

...to the extent that we cannot provide a certain kind of compositional semantics for some language...we will be unable to straightforwardly express in our terms the *fine-grained intentional content* of some of the assertions made in that language...some of their *intentional states* may not be fully expressible in our own language (Hirsch, 2005: 80 [emphasis added]).

What Hirsch is suggesting here, in light of what has been discussed, is that although the sentences ‘there are simples arranged table-wise^N’ and ‘there is a

half full. Thus from that point of view, we can prise apart the difference between the glass being ‘half-full’ and ‘half-empty’ at the level of intensions. However, we can imagine a different world where the glass magically comes into existence at t₁, already containing beer to 50% of its capacity. The point here then is that *this* glass is equally truly describable as being ‘half-full’ or ‘half-empty’, but that these are still nonetheless different ways of conceptualising the same fact that there is a glass containing beer to 50% of its capacity.

table^{CSO}', in the same context of utterance, will each adequately and truly describe the same fact, each sentence is a different way of *conceptualising* or *representing* that fact. We can make more sense of this idea by thinking about Jespersen's idea again. Although the sentences 'the glass is half-full' and 'the glass is half-empty' are perfectly adequate ways of describing the same fact, they are not equally good at describing someone's way of conceptualising that fact. Similarly, for Hirsch, whilst 'there are simples arranged table-wise^N' and 'there is a table^{CSO}' are equally good at describing the same fact, they are not equally good at describing someone's way of conceptualising that fact. This then makes sense of the idea that the nihilist will be unable to straightforwardly express the *intentional state* of a CSO-speaker when they assert 'there is a table^{CSO}' in their own language – the *concepts* the CSO-speaker is using to express that fact are ones that are not available in the N-language. That this is the sort of thing Hirsch has in mind is further demonstrated by the following quotation:

the basic idea of quantifier variance can be nicely formulated by saying that the same (unstructured) facts can be expressed using the different concepts of "the existence of a thing", that statements involving different kinds of quantifiers can be equally true by virtue of the same (unstructured) facts in the world...I can say in my A-language that the truth of an M-statement depends on the structured facts, even though what I mean by "the structure of the facts" is not what a speaker of the M-language means. (Hirsch, 2002: 59)⁵⁸

Here Hirsch is rather loose with his talk of structured and un-structured facts. But what is going on, it seems to me, is that Hirsch is suggesting that the same unstructured fact in reality can be expressed using different concepts of existence⁵⁹ – and if two sentences use different concepts to express the same fact we might say that two different 'structured facts' can express the same unstructured fact. Here, it appears that a structured fact is a purely a linguistic/mental notion and is not meant to imply any difference in structure in reality itself. In the above quotation then, Hirsch says that by using different concepts of existence, we arrive at two different descriptions of the same unstructured fact in reality. In the example we have been concerned with here, it also quite clear that 'table' and 'simples arranged

⁵⁸ The idea that the 'structure of the facts' can have different meanings in different languages will be important when we consider a potential solution to the criticism I am going produce in the next section.

⁵⁹ This issue will be of central importance in chapter eight.

table-wise' are in some sense conceptually distinct, for the same unstructured fact in reality can be conceptualised as a single object or as many different objects standing in relation to each other. The idea then is that whilst the fact in reality stays the same, we 'carve up' the fact in thought and language using the different concepts we have at our disposal. However, this does not then take away from the fact that each sentence is an adequate and true description of the unstructured fact itself.

We can liken this to a passage from Austin:

We are absolutely free to appoint *any* symbol to describe *any* type of situation, so far as merely being true goes...There is no need whatsoever for the words used in making a true statement to 'mirror' in any way, however indirect, any feature whatsoever of the situation or event; a statement no more needs, in order to be true, to reproduce the 'multiplicity', say, or the 'structure' or 'form' of the reality, than a word needs to be echoic of writing pictographic (Austin 1950; cited in Blackburn & Simmons 1999: 155)

What Austin is suggesting here then is that insofar as a sentence is true *of* a particular fact, its syntactic structure need not mirror the structure of the fact it is true of. The idea is that we are free in thought and language to carve up the fact itself as we see fit – we are free to appoint any symbol to describe any situation, as Austin puts it. This is our power as language users. In a similar way, Hirsch would argue that we are free to assign any character to any syntactic structure⁶⁰. We then have a sentence which will be true of some fact in reality⁶¹, but it need not be the case that the structure of the sentence should somehow mirror the structure of the fact in reality itself.

Indeed, Hirsch even goes so far as to say that we should give up on the notion of 'language shaped facts that are in the world independent of language' (Hirsch 2002: 59). What this suggests is that although Austin appears to leave room for the possibility that the facts themselves have an inherent, language independent structure, this kind of talk, for Hirsch, is nonsensical – we can only talk of the

⁶⁰ See chapter two, section 4.1

⁶¹ Of course it is possible to assign a character to a syntactic structure such that the resultant sentence is necessarily false – i.e. there is no fact in reality of which it is true.

structure of a fact from *within* some linguistic scheme⁶². This confirms the idea that for Hirsch the idea of a structured fact can only be thought of linguistically or mentally – the fact in reality itself is unstructured, and in virtue of the different concepts we use, we read structure into the fact in our thought and language. What this then seems to suggest is that Hirsch would favour an account of hyperintensional content which meant that it was merely representational. Thus he would explain the difference in hyperintensional content between the sentences ‘there are simples arranged table-wise^N’ and ‘there is a table^{CSO}’ as being merely representational.

Thus far then I have shown that Hirsch would favour this account of hyperintensionality. It seems to me however, that his position absolutely *requires* that hyperintensional differences in content have to be merely representational and not arising from the structure of the world itself. If the difference in hyperintensional content between the sentences ‘there are simples arranged table-wise^{CSO}’ and ‘there is a table^{CSO}’ is *not* representational, then it can only be that the difference in content is expressing structural features of reality that cannot be expressed merely at the level of intensions. If that is the case, then it looks like the CSO-language has the resources to track structural features of reality that the N-language does not possess. That is to say, if ‘there are simples arranged table-wise’ and ‘there is a table’ are not just different ways of representing the same fact, but do indeed express different but necessarily co-existent facts, then it is the case that the N-language (which does not refer to tables) is not sufficiently expressive enough to distinguish between distinct but necessarily co-existing facts whereas the CSO-language is. Thus saying that the hyperintensional difference between these sentences ‘there are simples arranged table-wise^{CSO}’ and ‘there is a table^{CSO}’ is due to the fact that they express necessarily co-existing but distinct facts is to *actually take sides in the first order debate*. Therefore, in order for the debate to be verbal, the difference has to be merely representational. That way, the nihilist can charitably make sense of the apparent difference in fine-grained content between his own sentence ‘there are simples arranged table-wise^N’ and the CSO sentence ‘there is a table^{CSO}’ – he can regard this difference as being merely different ways of representing the same fact in reality. Thus the official explanation of any

⁶² This idea that it does not make sense to talk of the structure of the facts outside of some language will be of central importance in chapter eight. It might even be apparent to the reader at this stage that to claim that we should give up on ‘language shaped facts’ in the world could be potentially troubling for Hirsch.

hyperintensional difference in content must be that it is due merely to representation. However, this creates a problem for Hirsch.

5: The CSO speaker does not regard the hyperintensional differences to be merely representational

In the previous section it was argued that Hirsch needs it to be the case that the hyperintensional distinction in content between the CSO sentences ‘there is a table’ and ‘there are simples arranged table-wise’ is merely due to different representations of the same fact. We can represent this as a meta-level sentence in the following way:

(H) The sentences ‘there is a table’ and ‘there are simples arranged table-wise’ are true by virtue of the same fact.

The CSO speaker will actually *deny* H, and will argue instead that ‘H2’ is true:

(H2) ‘The sentences ‘there is a table’ and ‘there are simples arranged table-wise’ are true by virtue of distinct facts’.

Why will the CSO deny ‘H’ and instead assert ‘H2’? Quite simply because they do not take their own object-level sentences ‘there are simples arranged table-wise’ and ‘there is a table’ to be true of the same fact in reality. They think these sentences are made true by distinct but necessarily but co-obtaining facts. Let’s say that ‘H’ and ‘H2’ are examples of H-sentences. These sentences either assert that two object-level level sentences are true by virtue of distinct facts (as in ‘H2’) or they are true by virtue of the same fact. (as in ‘H’). In order for H-sentences of the form of ‘H2’ to be interesting, it will have to be the case that the object-level sentences being talked about are intensionally equivalent in the object language, otherwise the H-sentence will be trivially true. For example, ‘the sentences ‘there is a table’ and ‘there is a bridge’ are true by virtue of distinct facts’ is a perfectly good H-sentence, albeit one that is obvious and uninteresting. ‘H2’ is interesting because the object-level sentences under consideration are intensionally equivalent. The idea is that if ‘H2’ is true, it tells us that there are certain facts in reality that cannot be discerned merely at the level of intensions.

It is quite clear then that the CSO-speaker is committed to the meta-level sentence 'H2', since the facts expressed by 'there is a table' and 'there are simples arranged table-wise' cannot be discerned at the level of intensions. It is obvious the CSO-speaker will take 'there are simples arranged table-wise' and 'there is a table' as intensionally equivalent, because they also accept the relevant linking principle 'if there are simples arranged table-wise then there is a table' as necessarily true⁶³, but they also have the additional belief that the two sentences are made true by distinct facts, which is what is captured in 'H2'.

The problem with this is that Hirsch needs 'H' to be true in order for the nihilist to be able to charitably interpret the CSO as speaking the truth in his own language. If 'H2' is true, the charitable interpretation is not available to the nihilist. This is because the nihilist will not want to admit that the object-level sentences 'there is a table' and 'there are simples arranged table-wise' are both true in the CSO-language if they are true by virtue of distinct facts, since this would entail that the nihilist ontology is wrong. Thus, the nihilist could only charitably regard 'there is a table' as true in the CSO language if it expressed the same fact as what is expressed by the nihilist sentence 'there are simples arranged-table wise'. If the meta-level sentence 'H' is true, then that charitable interpretation is open to the nihilist. If it is false however, then the interpretation is not available.

Remember that a nihilist could 'fake' it as a CSO by speaking CSO-alt, which was phonetically and intensionally equivalent to the CSO language, except the nihilist took the linking principles of CSO-alt to be analytic⁶⁴. However, quite clearly the nihilist would be caught out in his fakery if asked what meta-level sentences he thought were true (again, see chapter three for more on this). It is the same once again in this instance. The nihilist can fake it as a CSO-speaker by speaking CSO-alt, and accepting 'H' as true. In this way he can assert phonetically and intensionally equivalent object-level sentences to the CSO speaker, all the while by not lying or betraying his true ontological commitments. When the nihilist asserts 'there is a table' whilst speaking CSO-alt, by his secret commitment to 'H' he takes 'there is a table' to express no additional ontological commitment to his original nihilist sentence 'there are simples arranged table-wise', understanding the sentence 'there is a table' to be nothing but a different representation of the same fact.

⁶³ But they do not take this sentence to be analytic of course – see chapters three and four.

⁶⁴ See chapter three for this.

But once again the nihilist will be caught out in his fakery if he is asked whether he assents to 'H2'. He cannot assent to 'H2' without lying or betraying his own ontological commitments. This shows once more that a difference in ontological positions can be masked at the object-level and can only be teased out by attending to what meta-level sentences one accepts. It shows once again that two speakers can be in agreement in terms of the object-level sentences they assert whilst differing over their ontological commitments (as the CSO speaker and the nihilistic CSO-alt speaker demonstrate this). Consider that there could be two varieties of nihilist, one who spoke a straight up nihilist language, such that in that language 'there is a table' is false, whilst another nihilist spoke CSO-alt, accepting 'H' as true, and thus accepting that 'there is a table' is true whilst understanding that as being merely a different way of representing the fact that there are simples arranged table-wise. Thus it is quite obvious that the dispute between the nihilist and the CSO-alt speaker is merely verbal, but it is *not* obvious that the CSO language is identical to the CSO-alt language, since the CSO speaker explicitly denies the truth of the meta-level sentence 'H'. Thus Hirsch simply cannot assume that the CSO language is identical with the CSO-alt language without begging the question against the CSO speaker. If he does make that identification anyway, then is not taking the debate on his own terms and ends up analysing a debate that no one is actually having. Once again, it is clear that the dispute between the nihilist and the CSO-alt speaker can be verbal, but that is not the same as the dispute between the nihilist and the CSO speaker. Hirsch cannot merely *insist* that this is the case – bare insistence is not an argument – Hirsch must *show* that there is no difference between the CSO language and the CSO-alt language.

6: Potential solutions for Hirsch

It now becomes clear that the problem of hyperintensionality for Hirsch is an analogue of the analytic problem. It will also be seen that there are a variety of solution to this new problem which are also analogues of the variety of solutions to the analytic problem, and they are as equally problematic for Hirsch also.

One might wonder that if the notion of hyperintensional distinctions in content is causing a problem for Hirsch, then why doesn't he just simply do away with the notion of hyperintensional distinctions in content and argue that there is only one grade of content – coarse-grained content that is captured at the level of the intention of a sentence. This first potential solution then is to argue that there is no

difference in fine-grained content after all. This clearly will not work for if there is no fine-grained distinction in content then the sentences ‘there are simples arranged table-wise^{CSO}’ and ‘there is a table^{CSO}’ say exactly the same thing. This is worse than saying that they are different representations of the same fact in reality because now they are not even different representations of the same fact in reality - they are just syntactically different in a way that has no effect on meaning whatsoever. Insofar as the CSO speaker *does* regard them as differing in meaning and indeed to express distinct facts in reality, this is not really a tenable solution to the problem for Hirsch, and merely bulldozes over an important distinction in order to force his merely-verbal claim through.

Another potential solution is to argue that there isn’t really a difference between saying that hyperintensional distinction in content are grounded in different representations of the same fact or that they are due to fine grained distinctions in the world itself. As we saw in chapter three, Hirsch could have made a similar argument against necessarily true sentences which were analytic/synthetic – argue that there is no difference, or that the fact of the matter was simply impossible to determine. Those solutions were seen to be unacceptable to Hirsch. The first was overly reliant on verificationist style arguments, and since Hirsch explicitly denied he was a verificationist, that would force him to concede an important commitment he made. Equally, arguing that whether a sentence was analytic/synthetic was impossible to resolve just meant that the ontological debate was impossible to resolve, not that it was merely verbal debate. This again was unacceptable to Hirsch, according to his own commitments.

Thus what happens if Hirsch argues either that i) there is no difference between saying hyperintensional distinctions in content are due to the world itself or due to our representations, or ii) if he says it is simply impossible to determine this? If Hirsch opts for (i), then once again he is relying on a verificationist style argument to make it his point. He also collapses the distinction between what is actually out there in the world and what is merely a representation – since now there is no ‘real’ difference between a genuine feature of reality and a mere representation of that reality, it looks quite difficult now to explain the nature of the relationship between our representations of reality and reality itself. Even worse, it looks like by making this claim, Hirsch also gives up on the notion of being a robust realist – something he explicitly doesn’t want to do. If there is no clear difference between a mere

representation and a genuine feature of reality, then the claim that our representations do not determine reality looks difficult to articulate and maintain. Thus opting for (i) would quickly mire Hirsch in an obscure position, but he most certainly couldn't maintain his claims of being a robust realist or an anti-verificationist. Equally, if he opted for (ii), then this would simply mean that ontological debates were hard or impossible to solve – not that they were merely verbal. So both these potential solutions are unattractive for Hirsch.

The final solution is to make the dispute over 'H2' itself merely-verbal. Just as with the solution to the analytic problem discussed in chapter four, we allow each side not only its own object language, but its own attendant meta-language also, and each side's meta-level claims come out true in the relevant meta-language, making the dispute over 'H2' itself merely verbal. Thus 'H2' would be true in the CSO meta-language but false in the nihilist meta-language.

This initially appears to be a neat solution for Hirsch but it again involves the claim that nothing makes 'H2' itself necessarily true other than a linguistic decision in the CSO meta-language. Equally, 'H' would only be necessarily true in the nihilist meta-language owing to nothing more than a linguistic decision. In other words, the relevant H-sentences would have to be analytic. This again creates a problem for Hirsch's commitment to robust realism. Firstly, consider that as a realist Hirsch would need to be fine with the idea that our linguistic decisions do not determine what facts there are, but that our linguistic decisions determine what sentences are true of what facts. This seems like a fair formulation of realism – the facts are out there, and do not depend for their existence on our thought and language.

However, now it looks like our linguistic decisions *do* determine what facts there are, since in the CSO language, two intensionally equivalent sentences are true by virtue of distinct facts, whereas according to the nihilist those same sentences are true by virtue of the *same* fact, and both those claims are true! Thus it looks like our linguistic decisions at the meta-level are determining what facts there are, at least numerically speaking (since the CSO meta-language recognises two facts where the nihilist meta-language recognises only one). If the facts are 'out there' it shouldn't be possible to determine via linguistic decision how many facts there are. Against this, Hirsch can only argue that the fact-talk sentences like 'H' and 'H2' are analytic and thus do not tell us anything about the world – they do not *really* express anything about extra-linguistic facts at all. Thus Hirsch will have to say that all fact-talk at the level of the meta-languages is itself true by nothing other

than linguistic decision. This again makes it hard for him to maintain realism in any robust sense, since a robust realist will want to be able to genuinely speak about what facts there are – the notion of a fact should have some substantive purchase for a realist. But there is no language now in which the notion of a fact does have substantive purchase. Thus whilst there could be a ‘realist language’ in which the claim ‘linguistic decisions do not determine what facts there are’ is true, this is only analytically true, as there could be an anti-realist language in which the sentence ‘linguistic decisions determine what facts there are’ is true and that sentence would be analytically true in the anti-realist language. It again seems like a *robust* realist couldn’t allow for the possibility that the very question of realism itself was nothing but a verbal dispute – that to me seems to be conceding realism in any robust sense.

Thus we can see that the H-problem has presented us with another route to showing how Hirsch ultimately needs to commit to T-an (i.e. that all a priori necessary truths are analytic). What this shows is that Hirsch’s modified light form of Carnapianism generates two specific problems – the analytic problem and the H-problem, and that any viable solutions to these problems which allow the central claim that the debates in physical object ontology are merely verbal to be maintained will necessitate giving up on several commitments that distinguish Hirsch’s position from a thoroughgoing Carnapianism. A thoroughgoing Carnapianism will not suffer from the analytic problem or the H-problem, but that is because a thoroughgoing Carnapianism is rooted in verificationism and anti-realism. These two things Hirsch sought to avoid, but it is increasingly looking like he cannot avoid these notions if we want to maintain his central claim. Thus again we can see that Hirsch’s modified light version of Carnapianism is untenable.

7: Summary of chapter six

In this chapter I have argued that there is a new problem that Hirsch’s position faces to do with hyperintensional content, which I have called here the H-problem. That problem arises initially because Hirsch himself recognises a finer-grained distinction in content between intensionally equivalent sentences. The problem is that Hirsch owes us an account of this distinction in content. I argued that there could be only two accounts, and both are problematic for Hirsch. The first is that the distinction in content is due to the world itself – intensionally equivalent sentences can pick out necessarily co-obtaining but distinct facts, and thus we

require the use of hyperintensional operators to prise apart these facts. This explanation is not acceptable for Hirsch because it implies that the N-language does not have the ability to distinguish between necessarily co-obtaining facts whereas the CSO-language does, in which case the CSO-language is clearly a better language to use to describe reality. Thus, this account of the distinction in content would mean that the debate between the nihilist and the CSO was not verbal after all.

The other account of the distinction in content was that it was merely representational. This accommodates the nihilist, but not the CSO, because the CSO will regard some intensionally equivalent sentences of their own language as expressing distinct facts, and thus on this account of content the debate again would not be verbal. This then is the H-problem.

I then argued that Hirsch could resolve the H-problem in a similar manner to the way he might resolve the analytic problem. The problem with that solution however was that Hirsch once again is seen to be retreating into a thoroughgoing Carnapianism which is *not* the position he intends to end up in when making his central claim that the disputes in physical object ontology are merely verbal. Ultimately then what has been demonstrated thus far is that Hirsch must either give up his central claim, or he must accept that his modified approach is untenable – he must embrace thoroughgoing Carnapianism.

Chapter seven: The Sider/Hirsch debate

1: Introduction

In the previous four chapters the overall criticism of Hirsch that emerged was that quantifier variance doesn't seem to allow us to say that debates in physical object ontology are verbal in Hirsch's intended sense because his meta-ontological position tends to infringe on the first order debates in non-neutral ways – specifically there are two problems of this variety – the analytic problem, which was discussed in chapter three, and the H-problem, which I developed in chapter six. It was also seen in the final section of chapter six that the solutions to these problems were similar, and that the proposed solutions also had similar problems. The most serious of these problems for Hirsch is that the proposed solutions drive him to a commitment to certain controversial assumptions from the meta-ontological perspective. The result of this is that proponents of serious ontology are given no reason to change their meta-ontological views because they can simply deny these controversial assumptions (see chapter six, section 7).

In this chapter, I want to now explore a different line of criticism against Hirsch's position. Recall that for Hirsch, accepting that there can be alternative meanings for quantifier expressions means accepting that there can be different 'semantic rules' governing sentence forms containing quantifier expressions (chapter two, section 2.2). Without accepting this possibility, quantifier variance, as a position, cannot get off the ground. As such then, Sider has developed a criticism of Hirsch arguing that the kinds of alternative languages that Hirsch envisions aren't possible owing the fact that there are *external constraints* on what quantifier expression can mean – the different semantic rules for quantifier expressions that Hirsch imagines are not eligible if there are external constraints on what quantifier expressions can mean.

This initial claim of Sider then – that quantifier variance is not even possible owing to external constraints on what quantifier expressions can mean, serves as the opening salvo in a lengthy exchange between Hirsch and Sider. The first phase of this debate can be roughly characterised as a meta-semantic dispute over what determines the meaning of our expressions – with Hirsch arguing that use can trump the sort of external constraints on meaning that Sider has in mind, thus allowing for the possibility of quantifier variance after all.

The second phase of the debate is characterised by the fact that Sider concedes the meta-semantic point that use *can* trump eligibility. However, he does not thereby concede the *meta-ontological* point that debates in physical object ontology are merely verbal. He argues that even if use trumps eligibility, we can make a decision to *allow* certain external facts to constrain what our expressions mean. These relevant facts, in the case of ontological disputes, will be what Sider calls the quantificational or ontological structure inherent in the world itself. Thus even if in ordinary discourse such facts do not constrain the meaning of quantifier expressions, we can make a decision to allow them to do so, thus we can have substantive disputes in ontology in a special language, *Ontologese*, in which external constraints *do* trump the use of quantifier expressions. Hirsch responds to this line of argument by arguing that any claim that a particular language *L*, is in some sense aligned to the world's structure, is a claim that will come out true in *L* itself, but not in any other language. This means that the question of which language is aligned to the structure of the world itself becomes *merely another verbal dispute*.

In this chapter I am going to explore both phases of the debate, and will argue that the second phase of the debate is plagued by the fact that Sider has in some sense misunderstood what Hirsch is claiming when he argues that questions about the world's quantificational structure are merely verbal. Sider, I argue, seems to understand this as a *metaphysical* claim to the effect that there is no objective quantificational structure to the world. This then is taken to inform Hirsch's meta-semantic claims that use trumps eligibility and his meta-ontological claim that debates in physical object ontology are merely verbal. This, I think, is not the right way at all to understand Hirsch's position. It is for this reason then that I will argue that Sider's criticism, whilst generating interesting points of contention, is ultimately not successful in challenging Hirsch.

In the final section of this chapter I will argue that this ultimately stems from Hirsch's commitment to neo-Fregeanism, a deflationary theory of reference, and a top-down approach to semantics. These commitments, I will argue, have a distinct implication for *how language relates to the world* – language primarily relates to the world at the level of whole sentences, not at the level of sub-sentential expressions. In not grasping this, Sider's criticisms do not work. However, in chapter eight I will then develop a criticism of Hirsch's position based on this very implication.

2: *Meaning is determined by use plus eligibility*

The thesis that meaning is determined by use plus eligibility arises from the fact that if meaning was solely determined by linguistic convention/use, then facts about how an expression is used within a community will not be enough to distinguish between equally good meaning-candidates for that expression, i.e. use alone does not give us semantic determinacy (Sider, 2001b: 190). Consider the expression ‘grue’, which is defined as following:

An object x is grue iff x is green and observed before t^{65} , or else is unobserved and is blue⁶⁶

A person’s use of the term ‘green’ up until some arbitrary time t is consistent with ‘green’ meaning *green* or *grue*. As such, if we imagine we are interpreting a linguistic community and have hitherto seen them referring to all and only green objects as ‘green’, we might reasonably conclude that ‘green’ means *green*, but we have no evidence to suggest that ‘green’ does not actually mean *grue*. It could be argued that we can wait until after t to see whether ‘green’ is used to apply to green or blue objects, thus settling whether ‘green’ means green or grue. However, that then does not rule out that ‘green’ means grue+ (where we simply take t to be a further time in the future). Additionally, we might ask someone in the linguistic community whether ‘green’ refers to all and only objects which are *this* colour (and then pointing at a green object). However, even then, if they say yes, this doesn’t settle the issue because their use of ‘colour’ could be taken to mean *schmolour*. If grue is a schmolour, then when we ask if ‘green’ means *this particular colour*, what we are asking is whether grue is this particular schmolour, which will be true⁶⁷. Thus with sufficient permutations in the meanings of other terms, it will be indeterminate on use alone whether ‘green’ means green or grue.

A response to this kind of worry about meaning is to argue that there are certain *external constraints*, called eligibility constraints, which rule out certain interpretations of the meanings of expressions. One of the most well-known accounts of this is the notion of property naturalness, where a candidate meaning for an expression is more eligible if it expresses a more natural property. The idea

⁶⁵ T is some arbitrary future time.

⁶⁶ The original formulation comes from Goodman’s *Fact Fiction and Forecast*, 1955.

⁶⁷ These sorts of problems, which are about following the rules of use for a particular expression, are from Kripke’s *Wittgenstein on Rules and Private Language*, 1982.

of a natural property itself is cashed out in terms of objective similarity between objects:

Among all the countless things and classes that there are, most are miscellaneous, gerrymandered, ill-demarcated. Only an elite minority are carved at the joints, so that their boundaries are established by objective sameness and difference in nature. Only these elite things and classes are eligible to serve as referents (Lewis, 1984: 227).

The idea then is that *grue* is not an eligible meaning candidate for the expression 'green' because there is a more natural property, *green*, which fits equally well with the use of the expression. Sider assumes this property-naturalness understanding of eligibility for the meaning of expressions and argues that meaning is determined by use plus eligibility (Sider 2001b: 191), which is to say that facts about how an expression is used (our meaning determining behaviour) and external facts about the world itself jointly determine the meaning of our expressions. From such a thesis, Sider considers three important outcomes:

- 1) Semantic indeterminacy is not ruled out by the thesis that meaning is determined by use and eligibility.
- 2) The *use* of an expression may be consistent between two different linguistic communities whilst that expression has different *meanings* in the respective communities.
- 3) Charitable truth-preserving interpretations of expressions can be ruled out by considerations of naturalness.

Let's consider first the claim that *semantic indeterminacy is not ruled out by the thesis that meaning is determined by use and eligibility*. Sider argues that if an expression *A* has two candidate meanings, then it might be the case that neither considerations of use nor eligibility can determine what *A* means. In such a case, the meaning of *A* is said to be indeterminate (Sider, 2001b: 189). Sider's example of such a case is the criteria of personal identity. He argues that the use of 'person' is such that it is indeterminate between whether we mean a bodily-continuous person or a psychologically-continuous person – sometimes we often speak of the same person as surviving a severe change in personality if the body remains intact, whilst on other occasions if there has been severe personality change and bodily continuity, we will lament the loss of a particular person – we will say they are no longer the same person (Sider, 2001b: 195-6). These considerations of how we use

the expression 'person', and thereby the persistence conditions we associate with the *concept* of a person, indicate that our use of 'person' is indeterminate between these meanings – Sider says that 'ordinary thought contains two concepts of persisting persons, each responsible for a separate set of intuitions, neither of which is *our* canonical conception to the exclusion of the other' (Sider, 2001b: 197).

Sider further argues that in the particular case of personal identity, there is no candidate meaning which expresses a more natural property – the notion of a bodily-continuous person and a psychologically-continuous person are equally eligible to be meant by our term 'person' (Sider, 2001b: 199). This, Sider argues, is only the case if both psychologically-continuous persons and bodily-continuous persons are acceptable entities in the ontology of both sides in the debate about what it is to be a person. Thus if one side of the debate doesn't accept the existence of psychologically-continuous persons, they will not accept the idea that the property of being a psychologically continuous person is an eligible candidate for the meaning of the expression 'person'. As Sider states, 'the no-fact-of-the-matter thesis is only *conditionally* established' (Sider, 2001b: 201). The point Sider is making here is that conceptual analysis can only take us so far – at some point we will bump up against issues of ontology which need to be settled before we can determine what our expressions mean. Thus it might turn out that the expression 'person' can *only* mean 'bodily-continuous person' if the true ontology doesn't contain any entities that could serve as a candidate meaning for 'psychologically continuous persons'. However, that of course supposes that questions of ontology themselves are not susceptible to being shown up as merely verbal disagreements, which is precisely what Hirsch says. Sider then needs to demonstrate that ontological disagreements are indeed substantive and cannot be 'analysed' away as nothing but verbal disagreements. To do this let us look at two other important considerations arising from the thesis that meaning is determined by use plus eligibility.

- 2) *The use of an expression may be consistent between two different linguistic communities whilst that expression has different meanings in the respective communities*

In making this claim Sider wants to say that two linguistic communities may possess the same concept, but external environmental differences mean that whilst the concept is the same across the different communities, the concepts denote

different kinds. This is basically Putnam's Twin Earth thought experiment. If we recall, Twin Earth contains a watery substance that is qualitatively identical to water and fills the lakes and seas of Twin Earth just as water does on Earth. The difference is that the chemical structure of Twin Earth water is XYZ and not H₂O. The implications of this, for Sider, are the following: Since Twin Earth water fulfils exactly the same role on Twin Earth as water does on Earth, then the use of the expression 'water' is the same across the two linguistic communities. For Sider, this is enough to say that the concept associated with the expression 'water' is shared between Earthlings and Twin Earthlings (Sider, 2001b: 191). However, the *meanings* of the expressions differ because the Earth-expression 'water' refers to H₂O in every possible world and the Twin-Earth expression 'water' refers to XYZ in every possible world. The implication of this is that if an Earthling said 'Twin Earth contains water' he would be speaking falsely, because we evaluate the truth of what he says in terms of what the words mean in his own language. As such, because in his language 'water' refers to H₂O, he is effectively saying that 'Twin Earth contains H₂O', which is false. The important consideration here then is that external facts that are unknown to a speaker will partly determine what their expression means⁶⁸.

3) *Charitable truth-preserving interpretations of expressions can be ruled out by considerations of naturalness*

Drawing from the previous consideration then, it seems that external constraints that are unknown to a speaker may rule out charitable truth-preserving interpretations of what they say. Sider thus describes two communities - our community, and an unenlightened community who have yet to discover the difference between gold and fool's gold. In such a case, Sider argues that this will result in a difference in how the expression 'gold' is used in the respective communities - we will discriminate between gold and fool's gold, whereas members of the unenlightened community will use the expression 'gold' to refer to both substances. Thus there is a clear difference in how we use the expressions, and this, for Sider, is enough to say that there are different concepts associated with

⁶⁸ Interestingly this doesn't seem to have anything to do with naturalness but rather the importance of the causal relations between speakers and objects in establishing reference. One might argue that the important lesson to draw here is that the descriptive conditions which fix how the term 'water' is to be used on Earth and Twin Earth are the same (if we make some assumptions about the scientific knowledge of the respective communities), and this settles how the expression 'water' is *used* on Earth and Twin Earth, but the descriptive conditions alone do not fully determine the *meaning* of the expressions because of the different 'hidden' properties of water and Twin-Earth water.

the expression 'gold' in the respective linguistic communities. However, he argues that because the property of *being gold* is more natural than the property of *being gold or fool's gold*, then eligibility constraints mean that in the unenlightened community, their expression 'gold' actually expresses the property of *being gold* (Sider, 2001b: 191). This has the following implication, which is only implicit in what Sider says. If there is some fool's gold in front of us, and a member of the unenlightened community says 'that is gold', we are faced with an issue of interpretation. Considerations of use alone would allow us to say that when the member of the unenlightened community says 'gold' they mean *gold-or-fool's gold*. As such, we would be charitably interpreting their sentence 'that is gold' as true. However, *considerations of eligibility rule out this possibility*. Thus even if everyone in the unenlightened community assents to the claim 'that is gold', considerations of eligibility overrule community-wide consensus. Their expression 'gold' means *gold*, just as it means *gold* in our community. The result of this is that the sentence is false. The important lesson here is that use alone cannot determine the meaning of an expression – even if an entire community is on board with how that expression is used. We must consider eligibility as well, and in this particular instance eligibility overrules considerations of use.

It is this final consideration that Sider uses to make his case that debates in ontology cannot be verbal. This is because eligibility constraints on the meaning of quantifier expressions rule out the charitable truth-preserving interpretations of quantifiers that Hirsch requires for his meta-ontological claims to work. In short, Sider denies that quantifier variance is possible.

Sider points that a deflationist such as Hirsch will argue that the meanings of quantificational expressions are exhausted by the rules of use of sentences containing those expressions (Sider, 2001b: 204). In other words, Hirsch argues that the *only thing* that determines what our quantificational expressions mean is use – there is no eligibility constraint in play. As Sider points out, on this sort of view

Quantificational expressions [...] get their meanings from the rules of language adopted by those that use them. The nihilist uses different rules for the quantifiers than do the chaste endurantist or the defender of temporal parts [...] In the nihilist's linguistic framework, there is no rule allowing one to infer ' $\exists x$ is made up of *a* and *b*' from

the assumption that a and b denote objects. This rule is included in the frameworks of the chaste endurantist and the defender of temporal parts [...] If the meanings of quantificational expressions are exhausted by rules of this sort, it might be argued that our meaning determining behaviour does not determinately settle which rules govern quantification (Sider, 2001b: 204).

As we can see then, the difference in meaning between the nihilist quantifier and what Sider calls the chaste endurantist⁶⁹ turns on whether one accepts 'If a and b denote objects then $\exists x: x$ is composed of a and b ' as rule of use in one's language.

Sider has a problem with this. He claims that 'intuitively, there is nothing the opponents of the nihilist can stipulate about the existential quantifier that will ensure that ' $\exists x$ is made up of a and b ' is true, *provided they use \exists as a quantifier*, for there simply may not *be* a third object other than those denoted by a and b ' (Sider, 2001b: 204) [original emphasis]. He further points out that a community is free to use a symbol or sentence to mean anything, but at a certain point it no longer makes sense to say that the sentence has anything to do with quantification or existence. In his example, he argues that a community can define ' $\exists x\phi(x)$ ' to mean that 'Nelson Goodman says that some object satisfies $\phi(x)$ ' (Sider, 2001b: 204). At this point, Sider says, ' \exists ' no longer has anything to do with *existence*' (Sider, 2001b: 204).

If we compare what Sider is saying here to what he says about the unenlightened community and their use of the expression 'gold', it suggests that even though the chaste endurantist community use ' \exists ' in such a way that 'If a and b denote objects then $\exists x: x$ is composed of a and b ' appears to be a rule of use for ' \exists ', this is *no guarantee* that there is anything that satisfies x in ' $\exists x: x$ is made up of a and b ' even if a and b do denote. Thus it may still be false, despite a community-wide consensus that ' $\exists x: x$ is made up of a and b ' is true. This is similar to the unenlightened community who assert 'that is gold' in the presence of fool's gold. Despite a community wide consensus, it is still possible that external constraints about what an expression means makes the assertion false.

⁶⁹ Understood to be a position where certain composite objects are admitted but not all the composite objects of common sense.

2.1: External constraints on the meanings of quantifier expressions

The question now is what external constraint could there be on the meaning of the quantifier? Sider argues that ‘just as the world comes ‘ready-made’ with natural properties and relations, it also comes ready-made with a domain of objects. This domain is extremely eligible to be meant by quantificational expressions’ (Sider, 2001b: 205). Thus, in order for any sentence of the form $\exists x\phi(x)$ to be true, there has to be some element in the domain which can satisfy x . We cannot simply infer $\exists x\phi(x)$ from the truth of some other sentence – this does not guarantee that $\exists x\phi(x)$ is true. Sider here is drawing on a fairly intuitive implication – there are some domains of quantification where we cannot stipulate what elements are in that domain. If two people are talking about the same domain of objects, it is a fact about whether there is some element in the domain that can settle their debate – it can’t be that for one person there is an element in the domain and for another person there isn’t. For example, take Fermat’s last theorem in question form – are there three positive integers, a , b and c which satisfy the equation $a^n + b^n = c^n$ where n is an integer greater than 2? The domain of positive integers is not up for grabs in this debate – if someone claims ‘there are positive integers which satisfy Fermat’s last theorem’, this is true only if there are three elements in the set of positive integers which do indeed satisfy the theorem. We can’t simply add inferential rules to our quantificational expressions to make this come out true in one person’s mouth and false in another *if both are talking about the same domain of objects*. There is a domain of positive integers, and whether we like it or not, it either contains elements which satisfy Fermat’s last theorem or it does not. Similarly, Sider is arguing that if everyone is in agreement that we are talking about the domain of every (concrete) thing that exists, either there is some element in that domain that satisfies $\exists x\phi(x)$, or there isn’t. This is why Sider says ‘it is hard to see why the different rules of inference should be regarded as alternate meanings for the quantifier, rather than alternative *beliefs* about what exists’ (Sider, 2001b: 204).

What can Sider mean by the claim that the domain is extremely eligible to be meant by quantifier expressions? Imagine I am a CSO and my opponent is a universalist, and we both claim to be talking about the domain of concrete objects. The most eligible interpretation of our quantifier expressions then is that they simply quantify over objects in that domain. However, the universalist will say that ‘there is something composed of Clinton’s nose and the Eiffel Tower’. According to

Hirsch, this claim is true providing we *use* the universalist concept of existence, which is captured by the rule ‘if a and b denote then the sentence ‘there is something composed of a and b is true’. As such then, from the CSO perspective, given what we believe about the objects in the domain of concrete objects, in order to charitably interpret the universalist claim ‘there is something composed of Clinton’s nose and the Eiffel Tower’ as true, we have to interpret that there are additional rules governing how the quantifier is used in the universalist language that go beyond simply quantifying over what exists, because from our own perspective, there simply isn’t anything in the domain to stand as the value of x in ‘there is some x such that x is composed of Clinton’s nose and the Eiffel Tower’, thus whatever is going on the universalist language when they assert this must be more than a basic act of quantification. Thus from the perspective of the CSO, the charitable truth-preserving interpretation of the universalist quantifier is that it has associated with it additional inferential functions⁷⁰ that go beyond merely quantifying over the objects in a stipulated domain.

The problem with this, for Hirsch, is that if there is an eligibility constraint on what quantifier expressions can mean, then this charitable truth-preserving interpretation is ruled out, and as such the CSO must instead interpret the universalist assertion as false on *the only available interpretation of quantifier expressions*. The idea here then is that even if the whole universalist community spoke such that their use of quantifier expressions was consistent with the rule ‘if a and b denote then the sentence ‘there is something composed of a and b ’ is true’, they would still be speaking falsely when they asserted ‘there is something composed of Clinton’s nose and the Eiffel tower’. This is similar to the way we interpret the unenlightened community’s assertions involving the expression ‘gold’. Just because their *use* of ‘gold’ is consistent with it meaning gold or fool’s gold, this truth-preserving charitable interpretation is ruled out because eligibility constraints tip the scales in favour of the expression ‘gold’ meaning gold, such that the assertion ‘that is gold’ in the presence of fool’s gold comes out false on the only available interpretation. Thus, from the CSO perspective, on the only available interpretation of quantifier expressions, the universalist assertion ‘there is something composed of Clinton’s nose and the Eiffel Tower’ is false – a truth-preserving interpretation is ruled out. In keeping with the ‘gold’ example, we will say that whilst the universalist has a different *concept* of existence, evidenced by

⁷⁰ Hirsch would refer to these as ‘semantic rules’.

the different rules of use for quantifier expressions in their language, those quantifier expressions have the same *meaning* as our quantifier expressions. The result of this is that it is incorrect to say that the universalist is speaking a different language at all; rather the universalist is speaking the same language as us and is *making a factual error about what exists*. This is because the universalist language is supposed to be individuated by the different meaning of the quantifier expression in that language, but since there cannot now be this different meaning, the means by which we individuate the universalist language has disappeared. Thus the claim that the universalist can speak a different language where their assertions come out true cannot be right. As Hirsch himself remarks, if Sider is correct in his position then ‘...The ontological disputes are not verbal because it is impossible for there to be the different languages I describe’ (Hirsch, 2005: 90).

To summarise Sider’s position then, first of all he argues that the meanings of expressions are determined by eligibility constraints as well as our meaning determining behaviour (use), and that sometimes eligibility constraints overrule use in determining the meaning of an expression. A result of this, when applied to quantifiers, is that the most eligible meaning for a quantifier is simply its function of quantifying over a domain of objects. However, in order to charitably interpret the assertions of another side as true in an ontological debate, we have to argue that there can be additional inferential rules for sentences involving quantification which make those sentences true, even if (from our own perspective) the purported object being quantified over is not in the domain of objects. These kinds of interpretations of the quantifier are ruled out as ineligible according to Sider, and as such we cannot charitably interpret the assertions of another side in an ontological debate as coming out true. A secondary claim of Sider’s is that the world contains a domain of ready-made concrete objects, and it is the objects in this domain that are being discussed when we are engaged in debates in physical object ontology, and the objects in that domain are ‘not up for grabs’ so to speak – they exist out there in the world and no amount of fiddling with rules of quantification will make claims about what objects are contained in that domain true if the objects simply are not in that domain. Let us now consider Hirsch’s response to this initial criticism of Sider’s.

3: Hirsch's reply to Sider

Hirsch responds to this initial criticism of Sider in two ways. First he challenges Sider's appeal to a domain of ready-made objects in the world in order to ground the notion of a most eligible meaning for quantifier expressions. Secondly he argues that eligibility constraints on what expressions can mean are often defeated by use – the idea that there is an all-out external constraint on what expressions can mean is a radical thesis that seems false. I will now look at both of these claims in more detail.

Hirsch is unconvinced by Sider's appeal to the notion of 'logical joints' in reality and particularly the notion of a domain of ready-made objects in order to explain the eligibility constraints on the meaning of the quantifier. In one sense he regards the appeal to logical joints as ineffectual, in that he regards it as a mere repetition of the original constraint on quantifier meanings in metaphorical form; 'The image is of joints or grooves in the world into which properly functioning quantifier-like expressions easily fit. In the absence of the constraint the idea of there being such joints is meaningless' (Hirsch, 2005: 92). Thus the appeal to logical joints in reality, as far as Hirsch is concerned, will not help to explain Sider's insistence on there being a constraint on the meaning of quantifier expressions.

Hirsch is more wary of Sider's appeal to a domain of ready-made objects which the unrestricted quantifier is supposed to range over. Hirsch argues that the meaning of a 'domain of ready-made objects' itself will be different in different languages:

I [...] also believe in "ready-made objects" in the following prosaic sense: There are numerous objects in the world - rocks, rivers, trees, apples, planets, electrons - whose existence do not depend in any way on the existence of language or thought. These objects typically do not have temporal parts or sums, and that too does not depend in any way on the existence of language or thought. If I had been speaking [4D]-English instead of plain English, I would have correctly said, "These objects have temporal parts and sums, and that does not depend in any way on the existence of language or thought." It's essential in this area of philosophy to avoid a gross but somehow tempting use-mention confusion (Hirsch, 2005: 93).

Hirsch here is reminding us that when speaking plain English, it is not true to say ‘our linguistic decisions determine what exists’, but it is right to say that (now speaking a meta-language) linguistic decisions determine the meaning of the expression ‘exist’ (and as such truth-conditions of the sentences where this expression is used)⁷¹. Thus, linguistic decisions will determine the meaning a cluster of related terms – ‘exist’, ‘reference’, ‘something’ and ‘object’. Hirsch often repeats the phrase ‘as goes existence, so goes reference’ (Hirsch, 2002: 55), and it has already been discussed how Hirsch is committed to a deflationary view of reference (see chapter two, section 3.1). The basic idea is that a shift in how ‘exists’ is used will also effect a shift in how ‘reference’ and ‘object’ are used in a language because of the interrelated nature of these expressions. The important result is that sentences like ‘my quantifier ranges over everything that exists’ has a different meaning in different languages, and so will the sentence ‘my quantifier ranges over the domain of ready-made concrete objects’. Similarly, if a sentence such as ‘tables exist’ is true, it will trivially follow that ‘the word ‘table’ refers to something’ is true, and it will also follow that ‘there is an object in the domain of ready-made objects that satisfies the predicate ‘...is a table’ is true. The idea then is that the question ‘are there *F*’s in the domain of ready-made objects?’ is just to ask ‘do *F*’s exist?’, and this is a question that has different answers in different languages. Consider what Hirsch says in the following quote:

The [ME]-speaker will, of course, make the platitudinous disquotational assertion, “if something exists it is referred to by the word ‘something’.” Given what they mean by ‘something’ this sentence is trivially true. We cannot therefore ask the RC-speakers, ‘Is there a semantic restriction on the RC-quantifiers?’ since *that question is merely another form of the question whether there exists such things as mereological sums and temporal parts, a question which has different answers in the different languages* (Hirsch, 2005: 77)[emphasis added].

Here we see the basic idea that the very question of which language is using the most natural meaning for the quantifier itself is simply *a restatement of the original question about what objects exists*, and thus is merely a verbal question which will have different answers depending on what language one is speaking.

⁷¹ Recall the discussion from chapter three, section 5.

In this sense then, Hirsch seems to imply that the appeal to a domain of ‘ready-made’ objects cannot objectively ground which among the various alternative languages (e.g. the N-language, the CSO-language etc.) makes use of the quantifier with the most natural meaning. The idea then is that the original debate about which objects exists simply resurfaces at a higher level about which language has the most natural quantifier, which itself is a verbal question, according to Hirsch.

The second response that Hirsch gives is that Sider’s constraint on quantificational expressions is an unprecedented move since it is construed as an indefeasible presumption which cannot ever be defeated by appeals to use. Hirsch argues that even Lewis’s understanding of eligibility constraints on predicates in terms of natural properties was a defeasible presumption that could be defeated by appeals to charity (Hirsch, 2005: 95). The idea is that Sider’s constraint is too harsh, especially if it ignores charity to retraction, to be plausible. Remember that charity to retraction is the idea that reasonable people will improve the accuracy of their beliefs in the face of additional evidence or argument, and that a refusal to retract would indicate a difference in meaning in one of the key terms in the debate (see chapter two, sections 6.1-6.2). Let us consider Hirsch’s example of the Melville community. The Melville community are a linguistic community who speak like Ishmael from *Moby Dick*. Hirsch’s point is that the Ishmael character was aware of the differences between whales and fish, but didn’t care - Ishmael proclaims ‘Be it known that, waiving all argument, I take the good old fashioned ground that the whale is a fish...This fundamental thing settled, the next point is, in what internal respect does the whale differ from other fish.’ (Hirsch, 2005: 94; note 54). Thus in the Melville community, the word ‘fish’ expresses the property of *being a fish or a whale*. Now, this is insisted upon in the Melville community even when a scientist describes the internal differences between whales and fish. The Melville community is not disposed to retract their claim that ‘whales are fish’ even in the face of what the scientist says. Thus we ask whether it is plausible that despite the evidence being available and understood, the Melville-speaker speaks falsely when they assert that ‘whales are fish’. Surely when we are interpreting what a speaker means, we need to take into account what they would say under more ideal epistemic circumstances.

Let us consider again Sider’s example of the unenlightened community. If we decide to teach them the chemical difference between gold and fool’s gold, such that after this lesson, the members of the unenlightened community no longer asserted ‘that

is gold' when pointing at fool's gold, it would seem to suggest that all along the unenlightened community meant *gold* when saying 'gold', rather than *gold or fool's gold*. We thus imagine that when we teach them the difference between gold and fool's gold, they have learnt something about the *substance* gold, and *that* substance is what they had in mind all along when talking about 'gold'. However, the Melville community is aware of the biological differences between whales and fish, but this knowledge plays no role in how they categorise things. This suggests that naturalness as an external constraint on what expressions mean can be defeated if a linguistic community explicitly denies that that is what they meant by a particular expression. If we are a radical interpreter, the eligibility constraints on predicates in terms of property naturalness can seem like a good guide if we have nothing else to go on, but it seems reasonable to suggest that if the linguistic community themselves do not categorise things along natural property lines then this seems like we are just ignoring important data on what their expressions mean. Consider that the seafaring folk of the Melville community are not necessarily concerned with the internal morphological differences between whales and fish – for them the external likeness is more important. The idea seems to be that similarity can in some respects be interest relative, and objects can be categorised according to these relative interests – perhaps not every community is interested in objective similarity, and are more concerned with categorising things according to the specific goals/interests of that community. To summarise then, Hirsch claims that 'People have the right to use the word 'fish' to mean, roughly, 'creature that lives in the water and has a fish-like appearance,' and charity would surely determine that the Melville-community uses 'fish' in that way (Hirsch, 2005: 94). For Hirsch, Sider's constraint is an absolutely indefeasible presumption that rules this kind of interpretation out, and Hirsch finds this implausible. He argues that 'there is nothing that we know about the nature of language or interpretation that prepares us for this kind of absolute constraint or that makes it seem credible' (Hirsch, 2005: 96).

In summary then, Hirsch's initial response to Sider amounts to the following: first of all an appeal to a domain of ready-made objects cannot ground the notion of an external constraint on quantifier meanings, because answers to the question of what objects there are in the domain of ready-made objects is just a restatement of the question of what exists, which has different answers in different languages. Furthermore, the very idea of an all-out constraint on what an expression can

mean is found to be implausible given that certain communities are not disposed to revise their claims in the face of seemingly relevant evidence. This suggests that eligibility can often be overruled by use in determining what an expression means in a particular community. Thus it remains an open possibility that use can determine the meaning of quantifier expressions, in opposition to what Sider thinks.

4: Assessing the debate thus far

The first phase of the debate is focused primarily on the meta-semantic issue of whether use trumps eligibility when it comes to quantifier expressions. A related point of contention is what sort of external facts could be said to constraint the meaning of quantifier expressions. Sider makes the meta-semantic claim that eligibility trumps use when it comes to determining the meaning of quantifier expressions. He then claims that there is an *objective domain of physical objects*, and that the function of the quantifier is to range over that domain – any additional semantic rule associated with the quantifier are illegitimate additions to its function. When considering Hirsch's reply to these claims, it is important to understand whether Hirsch is challenging the ontological claim that there is an objective domain of physical objects, or whether he is simply making a meta-semantic point. What I mean by this will become clear as we go along.

Hirsch argues that the appeal to a ready-made domain of objects cannot ground the notion of a most eligible candidate meaning for quantifier expressions, because the statement 'x is in the domain of ready-made objects' is equivalent to saying 'x exists', and if 'x exists' has different truth conditions in different languages, so too will the statement 'x is in the domain of ready-made objects'. The point Hirsch seems to make is that to argue that there is an objective notion of a domain of ready-made objects is simply to argue that there is an objective answer to the question of what objects exists, but since Hirsch argues the latter is merely a verbal question, the former will also turn out to be a verbal question. Here then, Hirsch is not explicitly denying, nor is he affirming that there is an objective domain of ready-made objects – thus *he is not making any sort metaphysical claim*. Rather, he is bypassing this issue and making the *meta-ontological* point that if a statement of the form 'x exists' is true in *L* in virtue of the character it has in that language, then it will trivially follow that 'x is in the domain of ready-made objects' is also true in *L*. Thus for Hirsch we can only talk about the objective domain of ready-made objects

from *within* a particular language. This all follows from Hirsch's commitment to neo-Fregeanism, a top-down approach to semantics and a deflationary theory of reference. The idea then is we should not (here at least – though see chapter eight) construe Hirsch as making any claims about reality itself and whether it contains an objective domain of objects. He is simply making claims about what sentences come out true in different languages.

However, he does in addition to this go on to make a *meta-semantic* claim that use trumps eligibility, and that any meta-semantics whereby there is an indefeasible presumption that eligibility always trumps use is simply too extreme a view. This is where some of the confusion that plagues the second half of the debate begins to emerge. For to argue that use trumps eligibility (which is precisely what Hirsch claims) seems to imply one of two claims, each of which involving distinctly *metaphysical claims*. First of all we can construe the claim that use trumps eligibility as the claim that though *there are* relevant facts which *could* serve to constrain the meanings of certain expressions (i.e. a metaphysical claim), they simply do not (a meta-semantic claim) – we do not, as language users, have to pay heed to these kinds of facts if we do not wish to do so. The second claim would be to *deny* that there are relevant facts which could serve eligibility constraints for the meanings of certain expressions. This appears to be a distinctly metaphysical claim.

Let's compare this, for illumination, to the notion of natural properties. Imagine someone says use trumps eligibility when determining the meaning of predicate expressions. What this suggests is that *either* there is a hierarchy of natural properties (a metaphysical claim), but that these sorts of facts do not determine what we mean by our predicate expressions (a meta-semantic claim), *or* there is no hierarchy of natural properties such that it can be used to fix a notion of eligibility for predicate expressions (a metaphysical claim). What the latter means, in other words, is that there are no external facts that make any groupings of objects better or worse than any others. Accepting the former then suggests that whilst there are relevant facts that *could* determine the meanings of our predicate expressions, they simply do not. In order for naturalness to play a part in determining the meanings of our expressions, it first would have to be a decision to *allow* such facts to have a bearing on what we mean by our predicate expressions. This is evidenced in the case of the Melville community. The Melville community are aware of the differences between whales and fish, and seem to accept the idea that there is an

objective taxonomic rank of organisms in nature, but they do not regard such facts as important in determining what they want to mean by the word ‘whale’. The point is, by saying that use trumps eligibility we are admitting that there are objective streaks in nature which we *could* use as constraints on meaning, but first we would have to make a decision to do so – thus eligibility constraints are only constraints *if we decide that they are*, and there is no reason why any such decision is binding, hence the Melville community. However, we could also deny that the notion of eligibility constraints are intelligible at all by denying that the taxonomic rank of organisms itself is objective – it is simply one way of grouping things along relative lines of interest, this again looks like a metaphysical claim. In this case, it would be better to say that there are *no such things as eligibility constraints* rather than saying that use can trump eligibility⁷².

As we now discuss the second phase of the debate, it is important to note that Sider appears to construe Hirsch as someone who denies that there are such things as eligibility constraints – and thus construes him as someone who is making a metaphysical claim that the world itself lacks any inherent quantificational/ontological structure. However, it will be shown that this not really what Hirsch is claiming – once again he bypasses this issue and makes the meta-ontological claim that any sentence of the form ‘ ϕ is part of the world’s quantificational/ontological structure’ will have different truth-conditions associated with it in different languages – hence the debate itself will still be verbal. Let us then consider the second phase of the debate by considering Sider’s reply to Hirsch.

5: Sider’s reply to Hirsch

In this section I am going to consider Sider’s 2009 reply to Hirsch. Siders’ reply can be reconstructed as three claims; one metaphysical, one meta-semantic, and one meta-ontological. I will explain each point in turn and along the way will discuss how I think Sider has incorrectly construed Hirsch’s position in his attempt to respond to it. The three points are the following then:

- 1) There are relevant facts which can serve as eligibility constraints for the meanings of quantifiers. This then is Sider’s metaphysical claim.

⁷² Though of course one not need be a global sceptic about eligibility constraints, since some terms (e.g. highly theoretical terms) might be governed by eligibility constrained whilst other terms might not.

- 2) Even if use does trump eligibility when determining the meanings of quantifier expressions in ordinary discourse, we stipulate our desire to speak a language where the meaning of the quantifier *is* constrained by the relevant external facts. This then is Sider's meta-semantic claim, which is dependent on the truth of the metaphysical claim.
- 3) In that special language (which Sider terms *Ontologese*), we can have substantive disagreements over matters of ontology. This is Sider's meta-ontological claim.

Let's now consider each claim in turn. First let's consider the metaphysical claim that there are relevant facts which can serve as eligibility constraints for the meanings of quantifiers. What sort of facts could serve as eligibility constraints for quantifiers? In his 2001b, as we have already seen, Sider suggested that the most natural role for an (unrestricted) quantifier was to range over a ready-made domain of objects, and that there couldn't be any strange inferential role associated with the quantifier – thus the relevant facts are simply facts about the elements in the ready-made domain of objects. In his 2009, Sider takes a slightly different route to explain what he means. He argues that just as objects sharing a natural property are objectively similar objects, *facts* involving the existence of something are all *objectively similar facts*. Fact similarity then seems to be part of the *objective structure of the world* along with property similarity, and similarity between facts seems like a good place to look for eligibility constraints on the meaning of the quantifier. Consider the following:

...the question is whether *quantifiers* carve at the joints. To answer we should look to similarity between *facts*, not similarity between particulars. When each of the following sentences is true:

Ted is sitting

John is sitting

we have similarity between the facts: between the fact that Ted is sitting and the fact that John is sitting. Now, in this case there happens to be a further similarity: a similarity between the particulars Ted and John. But not in other cases:

Ted is human

Ted is located in North America

Here there is but one particular, and so we have no similarity between (distinct) particulars; but [...] we do have similarity between the facts expressed by these sentences, in virtue of the recurrence of the particular Ted in each fact. And finally, now, consider:

Something is human

Something is located in North America

If the existential quantifier carves at the joints, we again have fact similarity. There is some genuine commonality between cases in which something is human and cases in which something is located in North America. Each is a case of *something* being a certain way, and that is a genuine similarity. (Sider, 2009: 405)

In this example Sider is asking us to accept the idea that sentences of the form ‘something is F’ are all similar in that they express *objectively similar kinds of fact*. There is something lacking from this account because we need to know what two objectively dissimilar facts would be expressed by the two sentences ‘something is F’ and ‘something is G’ such that the occurrence of the word ‘something’ in both instances therefore constitutes an ineligible (or unnatural) meaning for the word ‘something’. Let’s try and provide a brief overview of what that account would look like. Let’s say first of all that the most eligible meaning for ‘something’ is that sentences of the form ‘something is F’ can only express facts where *something* is in some way. If a sentence of the form ‘something is F’ can also be used to express facts which do not involve *something* being in some way, then the use of the quantifier expression is unnatural. What other kinds of facts might ‘something is F’ be used to express then? If we consider things from the perspective of the nihilist, the CSO-sentence ‘something is a table’ does not actually express a fact where *something* is in some way, but rather expresses the *same fact* that is expressed by ‘some things are arranged table-wise’. Thus, from the nihilist perspective it appears as if the CSO-sentence ‘something is a table’ can express a fact where some *things* are so-arranged, which appears to be a different *kind* of fact from one where *something* is in a certain way. Additionally, as a nihilist we note that the CSO also says ‘something is a simple’, which *does* express a fact where *something* is in a

certain way⁷³. Thus we see that the CSO-language has two sentences of the form ‘something is F’ which express two different kinds of facts. Thus, the conclusion is that their quantifier does not carve at the joints as it does not *sufficiently distinguish between facts of different kinds*. This would make sense of the notion of a particular language being *aligned to the quantificational/ontological structure of the world*. If it is the case that the structure of the world is the way the nihilist says it is, for example, then the CSO-language is not properly aligned to that structure for the reasons stated here.

Thus the relevant facts which can serve as eligibility constraints on the quantifier are to do with the syntactic structural similarities between facts and the sentences containing quantifiers that are used to express them. Any quantified sentence which is true *of* a particular fact in some language but does not match syntactically the structure of the fact is an instance whereby the meaning of the quantifier is non-natural, and an instance where the language in which that sentence is true is not aligned to the quantificational structure of the world.

Against this, Sider construes Hirsch as someone who will deny that there are any such relevant facts – he construes Hirsch’s position as being committed to the *metaphysical claim* that asserts that there is no quantificational/ontological structure in the world itself. For example, he claims that that ‘a ‘negative’ thesis such as quantifier variance itself is a claim about the extent of the world’s structure...Quantifier variance is ‘just more metaphysics’” (Sider, 2009: 419). In that regard he considers the thesis of quantifier variance to be as epistemologically problematic as any other first-order ontological position (Sider, 2009: 419). Sider’s meta-semantic and meta-ontological claims that follow are conditional upon the assumption that the world contains quantificational structure. He then considers quantifier variance to be the assumption that the world contains no quantificational structure, and presumably thinks that Hirsch’s further meta-semantic and meta-ontological claims are conditional upon *that* assumption. Since he considers either assumption to be as epistemologically problematic as the other, he proceeds with his argument thinking he is as justified in asserting that the world contains quantificational structure as Hirsch is in denying it. However, when considering Hirsch’s reply to Sider, it will be seen that Hirsch doesn’t really

⁷³ Speaking loosely of course – something being a simple does not mean something is in a particular way of being a simple, rather we should say something has the property of being a simple, or rather something *is* a simple.

construe his meta-ontological claim as being conditional upon any sort of metaphysical assumption about the ‘extent of the world’s structure’. For now though, let’s proceed with Sider’s argument and look at his second claim.

His second claim is that even if use does trump eligibility when determining the meanings of quantifier expressions in ordinary discourse, we stipulate our desire to speak a language where the meaning of the quantifier *is* constrained by the relevant external facts. As mentioned in section 4, if we argue that use trumps eligibility, one way of understanding this is that there might only be relevant facts which constrain meaning if we as language users *allow* such facts to constrain the meaning of our terms (recall the Melville community). To concede then that in ordinary discourse use will trump eligibility is not a major concession on Sider’s part providing that we can then stipulate our intention to speak a language, *Ontologese*, where relevant external facts do indeed constrain the meaning of quantifier expressions. Again this is conditional upon the metaphysical assumption that the world does indeed contain quantificational structure.

Let’s *give* the speakers of ordinary English ‘there exists’; let us henceforth conduct our debate using ‘ \exists ’. We hereby stipulate that ‘ \exists ’ is to express an austere relative of the ordinary English notion of existence. We hereby stipulate that although the meaning of ‘ \exists ’ is to obey the core inferential role of English quantifiers, ordinary, casual use of disputed sentences involving ‘there exists’ (such as ‘Tables exist’) are not to affect at all what we mean by ‘ \exists ’. We hereby stipulate that if there is a highly natural meaning that satisfies these constraints then that is what we mean by ‘ \exists ’. Perhaps the resulting ‘ \exists ’ has no synonyms in English. Fine – we hereby dub our new language *Ontologese* (Sider, 2009: 412).

This imaginary stipulation to speak *Ontologese* can be thought of as performing the following function: We consider how sentences of the form ‘there exists...’ which are controversial and disputed amongst ontologists, including ‘there exists...’ sentences in ordinary English, are used. By intending to speak *Ontologese*, the *use* of these sentences in their respective languages is not to affect the meaning of the quantifier (beyond the core inferential role). Thus we might say that although in the CSO-language, ‘there exist tables’ is generally used and accepted as true, this is not to

affect the meaning of ‘there exists’ in *Ontologese*. The same goes for any other disputed sentence, such as ‘there exists arbitrary mereological fusions’. As such, although these may be construed as rules of use for sentences containing ‘there exists’ in the various languages, *these do not count as rules of use* in *Ontologese*. Thus if we imagine the 4d-language sentence ‘there exists arbitrary mereological fusions’ as being a rule of use, then from the truth of ‘A exists’ and ‘B exists’ we can infer ‘there exists a fusion of A and B’. Such an inference is not licensed in *Ontologese* – if ‘there exists a fusion of A and B’ is true in *Ontologese*, it is not the result of an inference using rules of language, but is true in virtue of whether it is a fact in reality that *something is a fusion of A and B*.

This then leads on to the final meta-ontological claim that in *Ontologese* we can have substantive disputes over matters of ontology. The committed four-dimensionalist will say that ‘there exists a fusion of A and B’ is true in *Ontologese* because there is some fact in reality which is such that *something is a fusion of A and B*, but the CSO and the nihilist will deny this. This is what Sider wants. In *Ontologese*, substantive ontological debate becomes a possibility again because we are using the most natural meaning for quantifier expressions. This is why Sider’s concession to Hirsch that the possible languages envisioned by Hirsch might be possible after all is not a self-undermining concession. It turns out that if one wants to do genuine ontology, one cannot rest content with speaking the truth in one’s own language – one must aspire to speak the truth in the ontologically best language. Sider claims, for example, ‘the goal of inquiry is not merely to believe many true propositions and few false ones. It is to discern the structure of the world ... employers of worse languages are worse inquirers’ (Sider, 2009: 401). Let’s now consider Hirsch’ reply to Sider.

6: Hirsch’s second reply to Sider

In the previous section I argued that Sider appears to construe Hirsch’s meta-semantic and meta-ontological claims as being conditional upon a denial that there are any external facts in reality itself that could be used to serve as eligibility constraints for quantifier expressions. In this section I want to show that this is not really how Hirsch understands his position and that Sider has possibly misunderstood Hirsch. For Hirsch the issue now is not whether or not the world contains any quantificational structure which one language might be aligned to. Rather his point is that any disagreement over which language is *Ontologese* is

merely a ‘recapitulation of the traditional ontologist’s verbal debate about what exists’ (Hirsch, 2008: 521). What Hirsch does is to argue that there is no sense in which we can talk of the objective structure of the world outside of some language – we will have to do it from *within* a language that is already *interpreted*. This will have the implication that for any language L , it will be true in L to say that ‘this language, but no other, is aligned to the world’s quantificational structure’. The idea then is that the question of which language is the ‘ontologically best language’ will have different answers in different languages – this debate itself is merely another verbal debate. Hirsch does this without making any distinctly metaphysical claims about the extent of the world’s quantificational structure. Let’s now examine his reasoning.

Hirsch argues that if each side in the debate wants to accept the idea of Ontologese, they will have to accept certain principles for determining when a language is Ontologese. These principles will have to be stated as sentences, and these sentences will have to be true in each of the alternative languages that the ontologists are speaking when they enter the ontology room. The result of this, however, is that the principles themselves, being true in each language, make it true to say in each language that ‘this language, but no other is Ontologese’. Thus, by accepting the principles for speaking Ontologese as true in one’s own language, it trivially follows that one’s own language is Ontologese! The point then is that the question of which language is Ontologese is merely a verbal question as the answer depends on what language one is speaking prior to entering the ontology room (Hirsch, 2008: 521).

The strength of this argument rests on two claims. The first is the claim that accepting the notion of Ontologese entails accepting certain principles as true sentences of one’s own language. The second is that the principles themselves, as true sentences of a language, do then indeed trivially show that it is true to say in L ‘This language, and no other language, is Ontologese’.

These two claims will be examined momentarily. First however, let’s look at the supposed principles themselves as formulated by Hirsch:

- 1) A necessary and sufficient condition for an ontological language L to be aligned to the world’s quantificational structure is that any sentence in L of the form ‘there exists such-and-such’ is true iff part of the world’s

quantificational structure consists in the fact that there exists such-and-such.

- 2) There exists such-and-such iff part of the world's quantificational structure consists in the fact that there exists such-and-such (Hirsch, 2008: 522)

Why does Hirsch argue that the principles must take such a form? Consider the first principle. As Sider suggests in his 2009 (as explained in section 5 of this chapter), if a sentence is true of some fact, but it is not syntactically structured the way that the fact is structured⁷⁴, then that sentence is not the most natural or eligible way of expressing that fact. As such, we must make use of the idea of a sentence's syntactic structure matching the syntactic structure of the fact it expresses. This can be done disquotationally, thus for example; 'there exists a table in region R' expresses the fact that there exists a table in region R. Hirsch then expresses this as it being 'part of the world's quantificational structure' that there exists a table in region R. Essentially, what the first principle is saying then is that if a language is to be Ontologese, sentences expressing facts in that language can only be true if they syntactically match the structure of the facts they are expressing. This is to be contrasted with 'regular' languages where a sentence can be true of a particular fact *even if the syntactic structure of the sentence does not match the structure of the fact*. This seems like something Sider would thus far be happy to accept – what makes Ontologese special is that there are stricter rules about how a sentence can be true of a particular fact, whereas in other languages it is merely how a sentence is used that determines whether it is true of a particular fact or not.

Hirsch argues that the first principle assumes the second (Hirsch, 2008: 521). I do not see it this way – rather it seems to me that if the first principle is a necessary and sufficient condition for a language to be Ontologese, then a sentence of the form of the second principle must be necessarily true in *Ontologese*, not in every language. This is a point that Sider seems to pick up on, which I will discuss in the next section. However, for now let's consider Hirsch's first claim that in order to accept Ontologese, each side in a debate will have to accept the above principles as true sentences of their own language.

⁷⁴ Which is assumed as possible if the constraint on the meaning of the quantifier is relaxed.

Hirsch argues that guiding ontologists to conduct their debate in Ontologese will mean getting them to accept the above principles as true in their own language. I will quote his reasoning in full:

A central part of the explanation of Ontologese is the idea of quantificational structure, an idea that seems clearly to imply the principles [1 and 2]. Therefore, the ontologists, who we are now imagining to speak the different languages, are in effect instructed by Sider to accept [1 and 2]. It follows from Sider's account, therefore, that [1 and 2] are true in all the languages (Hirsch, 2008: 521).

This again makes the questionable assumption that *both* 1 and 2 must be true in every language, whereas we will see in the next section that Sider challenges this. For now then let's just move on to Hirsch's next point:

Accepting 1 and 2 as true in all the ontological languages entails that the following is true in each language:

- 3) This ontological language (the one currently being used), but no other ontological language, is aligned to the world's quantificational structure (Hirsch, 2008: 522)

The chain of reasoning to arrive at this conclusion is presumably left as an exercise for the reader in Hirsch's original article, so here I am going to go through the steps to show how the argument is supposed to be working. First, we are assuming that principles 1 and 2 are true in *each language*, as this is what is required to get the ontologists to accept the notion of Ontologese. Let's take the CSO-language then as our example.

- i) There exists a table (already true in CSO)
- ii) There exists a table is true iff part of the world's quantificational structure consists in the fact that there exists a table (an instance of principle 2 – assumed true in CSO for now)
- iii) Part of the world's quantificational structure consists in the fact that there exists a table (from i and ii)
- iv) For CSO to be aligned to the world's quantificational structure, the sentence 'there is a table' is true iff part of the world's quantificational structure consists in the fact that there exists a table (an instance of principle 1 – assumed true in CSO)

- v) CSO is aligned to the world's quantificational structure. (from i, iii and iv)

The conclusion therefore is that the CSO-language is Ontologese. A result of this then is that from the CSO perspective it is also true that no other language is Ontologese. The reasoning for this is as follows:

Consider the first principle again:

- 1) A necessary and sufficient condition for an ontological language L to be aligned to the world's quantificational structure is that any sentence in L of the form 'There exists such-and-such' is true iff part of the world's quantificational structure consists in the fact that there exists such-and-such.

It is important to remember that this is a *sentence of the CSO-language*. However, we can substitute L with 4D and we can replace the sentence 'there exists such-and-such' with a sentence ' ϕ ' which is only true in 4D. Thus, we have a CSO sentence which is *about* the 4d-language and which mentions a true sentence of the 4d-language. Now consider the second principle

- 2) There exists such-and-such iff part of the world's quantificational structure consists in the fact that there exists such-and-such

In this sentence schema we take ' ϕ ' and *use* it as a sentence of our own CSO-language. So we have:

ϕ iff part of the world's quantificational structure consists in the fact that ϕ

Now, ϕ is true in 4D *but crucially it is false* in CSO. Thus, because principle 2 is a bi-conditional, and ϕ is an antecedent which is false, then the consequent is false. So in the CSO-language, the following is true:

It is *not* the case that part of the world's quantificational structure consists in the fact that ϕ

Now let's consider principle 1 again in light of this. There is a bi-conditional in 1:

' ϕ ' is true iff part of the world's quantificational structure consists in the fact that ϕ

But we already know that ' ϕ ' is true in 4D, and we also know that the consequent of this bi-conditional is *false* in CSO. Thus overall the bi-conditional is false. The result of this is that ' ϕ ' is true in 4d despite the fact that it is not the case that part of the world's quantificational structure consists in the fact that ϕ . Thus the necessary and sufficient condition for counting 4D as a language which is aligned to the world's quantificational structure is not fulfilled. Hence it is true to say, in the CSO-language, that '4D is not aligned to the world's quantificational structure'. Thus Hirsch has his conclusion:

This ontological language (the one currently being used), but no other ontological language, is aligned to the world's quantificational structure. (Hirsch, 2008: 522)

To summarise Hirsch's counter to Sider then, Hirsch is arguing that the notion of there being an objective answer to which language is Ontologese is nonsensical, as the claim that 'this language but no other is Ontologese' comes out true in each language. This makes the question of which language is Ontologese merely a verbal issue as the answer is different in each language. This results from the fact that the necessary and sufficient conditions for a language to be Ontologese must be stated as sentences which must be true in each ontological language. The result of this is that the necessary and sufficient conditions for a language to be Ontologese are *trivially fulfilled within each language*. Although in the next section we will see Sider challenge Hirsch on the specific formulation of his principles, we can observe the following point. The CSO-language is an interpreted language – thus it is already true in the CSO-language that 'There is a table in region R', for example. From this, it also trivially follows that the following is also true in the CSO-language: 'The CSO sentence 'there is a table in region R' expresses the fact that there is a table in region R'. The point then is that in each language, if we ask 'what is the structure of the fact expressed by the true sentence 'S'?', then the answer will simply be a restatement of the true sentence itself. Thus if in order for a language to be aligned to the world's quantificational structure, sentences expressing facts in that

language must somehow mirror or match the structure of the facts themselves, it appears that from *within* that language, this condition is fulfilled. Thus the sentence ‘this language is aligned to the world’s quantificational structure’ will then be true in each language – as such, one could argue that for Hirsch, talk of quantificational structure, and of a language being aligned to this structure, is just additional verbiage which disguises the original verbal dispute over the question of what exists⁷⁵. In making this point, it is also worth noting that Hirsch does not make any distinctly metaphysical claims about the extent of the world’s structure – thus his meta-ontological claim that the question of whether a particular language is aligned to the world’s quantificational structure is *not* conditional upon any metaphysical claim about the lack of structure in the world itself. It is once again important to bear this point in mind when we consider Sider’s response to Hirsch in the next section.

7: Sider’s second reply

In his 2014, Sider seeks to show that Hirsch’s claim that in each language it is true to say ‘this language, but no other, is Ontologese’ can be challenged. He argues that it is in fact intelligible to suppose that, for example, in *L* it is true to say that ‘*L*’s quantifier carves at the joints’. The caveat in Sider’s reply is that this is only so if we accept the metaphysical assumption that the world does indeed contain quantificational structure. Indeed, he claims that if this assumption turns out to be false then something like Hirsch’s position might be true (Sider, 2014: 566). Sider contends that this assumption is ‘dialectically appropriate’ – I shall consider whether this is the case later on. However, for now, let us presume that it is least dialectically appropriate to assume that there is objective quantificational structure in the world. That being the case, Sider argues that it is possible that the quantifiers of the 4d-language, for example, might be the quantifiers that are aligned to this structure (and thus 4d would be Ontologese). Given this then, it will be possible to formulate true sentences in other languages (such as the CSO-language) which state that 4d’s quantifiers are aligned to the world’s quantificational structure.

Sider reconstructs Hirsch’s original argument as follows:

⁷⁵ This is similar to how Hirsch responds to the issue of reference – the question whether of whether a singular expression ‘*a*’ *really* refers is just another way of asking whether *a* really exists – and thus both questions have different answers in different languages. See chapter two, section 3.1

- 1) $\neg\exists x(x \text{ is scattered})$
- 2) $\exists x(x \text{ is scattered}) \text{ iff } S(\exists x(x \text{ is scattered}))$
- 3) So $\neg S(\exists x(x \text{ is scattered}))$
- 4) ' $\exists x(x \text{ is scattered})$ ' is true in 4Dese
- 5) If ' $\exists x(x \text{ is scattered})$ ' is true in 4Dese and 4Dese is aligned, then $S(\exists x(x \text{ is scattered}))$
- 6) So, 4Dese is not aligned

(Sider, 2014: 568)

The above are taken to be sentences of the CSO-language, and S is taken to be a sentential operator which abbreviates the phrase 'part of the world's quantificational structure consists in the fact that...'. The above argument is supposed to show that in the CSO-language it is true to say that 4D is not Ontologese (i.e. is not 'aligned' to the world's quantificational structure). The argument relies on the principles we encountered in section 6, which Sider renders in the following way:

P: For any language L , if ' $\exists xFx$ ' is true in L and L is aligned, then $S(\exists xFx)$

Q: $\exists xFx \text{ iff } S(\exists xFx)$

These principles, according to Hirsch, need to be accepted as true in all the ontological languages if we are to introduce the idea of Ontologese. Unfortunately for Sider, as we have seen, the principles, being true in every language, result in the above argument (see section 6 of this chapter). 1 is true in the CSO-language, and 2 is an instance of Q, hence the intermediary conclusion 3. 4 is true as a CSO sentence as it is a sentence *about* what sentences are true in the 4d-language. 5 is an instance of P. 6 we get from 3, 4 and 5.

Sider argues that this argument only works on a key cashing out of the phrase ‘part of the world’s quantificational structure consists in the fact that...’. He argues that this can spelled out in two ways:

‘ $S_1\phi$ ’ is true in L iff ϕ is true and L ’s quantifiers carve at the joints

‘ $S_2\phi$ ’ is true in L iff ϕ is true *some language* L^* and L^* ’s quantifiers carve at the joints

Recall that Sider here is dialectically assuming as true that 4d’s quantifier carves at the joints. Thus if S is read as S_1 then all claims of $S\phi$ in CSO are false (Sider, 2014: 571). This is because $S\phi$ can only be true in CSO if ϕ is true and CSO’s quantifiers carve at the joints, but since we are supposing that the CSO’s quantifier does not carve at the joints, no instance of $S\phi$ can be true in the CSO-language.

If we now consider premise 5 of the argument, overall the conditional is false because according to Sider both conjuncts of the antecedent are true (the first is a premise in the argument whereas the second is true because of Sider’s dialectically appropriate assumption). The consequent however, is false, because all instances of $S\phi$ are false in CSO if S means S_1 . Because of this we cannot use premise 5 to conclude 6.

Secondly, if S means S_2 , then premise 5 is true because both antecedent and consequent are true. This is because $S(\exists xFx)$ will be true in the CSO-language if ‘ $\exists xFx$ ’ is true in 4d, since we are assuming that 4d’s quantifiers carve at the joints. Thus the consequent of 5 will be true. Furthermore, the antecedent is simply stating that ‘ $\exists xFx$ ’ is true in 4d and that 4d’s quantifiers carve at the joints. Thus overall the conditional is true, but we cannot use this to conclude 6.

Let’s now discuss this argument. First of all, it will be recalled from section 6 that I queried Hirsch’s insistence that both principles had to be true in each ontological language. The controversial principle was the following:

‘ ϕ ’ is true iff part of the world’s quantificational structure consists in the fact that ϕ

Let’s take an instance of this from the argument above:

2) $\exists x(x \text{ is scattered})$ iff $S(\exists x(x \text{ is scattered}))$

Hirsch wants to use this to conclude that:

$$\neg S(\exists x (x \text{ is scattered}))$$

in the CSO-language. However, if S means S_2 then he cannot do this. Consider that whilst the antecedent will be false in the CSO-language, the consequent will be true if S means S_2 , and as such Hirsch cannot conclude $\neg S(\exists x (x \text{ is scattered}))$. What Sider appears to be getting at is saying that whilst a sentence, S , can be false in L , an appropriate sentential operator on the front of that sentence can render it true in L . There is nothing strange in this idea - for example, fictionalist operators can be used to render ordinarily false sentences about fictional characters true. What Sider appears to be arguing then is the following: If we assume there is quantificational structure, we can also assume that 4d's quantifiers might be aligned to this quantificational structure whilst the quantifiers of other language do not. Thus whilst 'there is a scattered object' will be false in the CSO-language, the sentence 'part of the world's quantificational structure consists in the fact that there is a scattered object' will be true in the CSO-language.

One issue with Sider's argument here is that he assumes that 4d's quantifiers carve at the joints. He thus wants to then demonstrate that it will be true in each language that '4d's quantifiers carve at the joints'. Thus if we look at premise 5 of the argument, one of conjuncts of that argument is '4d's quantifiers carve at the joints', which for the purpose of Sider's argument we are assuming is true in the CSO-language. The problem then is that Sider appears to use this chain of reasoning to *prove* that '4d's quantifiers carve at the joints' is true in the CSO-language. This looks rather a bit circular. We can also see how this assumption affects the cashing out of S as ' $S_1\phi$ ' is true in L iff ϕ is true and L 's quantifiers carve at the joints. According to Sider, if L is the CSO-language, then L 's quantifier do not carve the joints, and thus all sentences of the form $S\phi$ in the CSO-language are false. But this only works if we assume that it is not a verbal question that 4d's quantifier carves at the joints. According to Hirsch's analysis, it comes out as true in the CSO-language that the CSO-quantifier carves at the joints, and so the assumption that 4d's quantifiers carve at the joints, if it is a sentence of the CSO-language, will be false in the CSO-language. Thus Sider is trying to foist a reading of ' S ' onto Hirsch that Hirsch need not really accept. One of the main issues for

Sider's response then is that his dialectically appropriate assumption seems quite inappropriate.

I think a more serious issue for Sider is that his argument is conditional upon the truth of the metaphysical assumption that the world does indeed contain quantificational structure, and that it has a specific quantificational structure independent of any language. As mentioned before, he positions himself against Hirsch, assuming that Hirsch's position makes the metaphysical assumption that the world does not contain quantificational structure. Since he considers the latter to stand in no better stead epistemologically than the former, he argues that his assumption is dialectically appropriate – and thus his argument is then meant to provide a stable platform from which people inclined to take ontology seriously can conduct their debates. The problem with this, I think, is that Hirsch's position does not appear to be conditional upon any metaphysical assumptions about the extent of the world's quantificational structure. Rather, he makes the *meta-ontological* claim that the sentence 'this language is aligned to the world's quantificational structure' comes out true in each language – he is making a claim about sentences in each language and is not making any claim about the world itself. The claim that the sentence 'this language is aligned to the world's quantificational structure' comes out true in each language for Hirsch will be derived *only* from his commitment to a top-down approach to semantics (with the addition of a deflationary theory of reference and a commitment to neo-Fregeanism – see chapter two, sections 3 – 4.2). Thus I think it is a mistake for Sider to set up the debate in the way he has – and I think this is due to a failure to appreciate Hirsch's approach to language and the implications this has for Hirsch's meta-ontological claims. In the final section of this chapter then I shall assess this particular issue.

Consider also the point that when Hirsch is not doing meta-ontology he just asserts what comes out true according to common-sense, i.e., what is true in the CSO-language (see chapter four, section 6). Thus when he asserts, from the CSO-perspective, that 'part of the world's quantificational structure consists in the fact that there is a table' he can recognise this is being true in the CSO-language, whilst recognising that in some other language it will be false, but that is *not* the language he is speaking. Thus any latent metaphysical commitment that Hirsch appears to have is only from the perspective of what is true in the CSO-language. Thus Sider's second reply does not appear to be effective.

8: Assessing the debate

The initial disagreement between Sider and Hirsch was over the meta-semantic facts which determine the meanings of quantifier expressions. For Sider, it made sense to talk of eligibility constraints on quantifier expressions, whereas for Hirsch it appeared that use alone could determine the meaning of quantifier expressions. An important notion for Sider is the idea that individual components of language – quantifiers, predicates etc., can be *substantively* related to the world in the relationship of carving at the joints of the world. It is this that allowed him to then argue that although two sentences could be true of the same fact, only one really got at the structure of the fact.

However, a proper understanding of Hirsch's views on language and how it relates to the world will reveal a number of interesting points. First of all, if we recall from chapter two, sections 4-4.1, it is not the use of individual terms that determine their meaning for Hirsch. Rather, how a *whole sentence* is used within a linguistic community determines its character (which determines the intension of the sentence in a context of utterance). Thus the relevant meta-semantic facts for Hirsch are how whole sentences are used, not individual expressions. Secondly, we should then remember Hirsch's commitment to neo-Fregeanism and a deflationary theory of reference. For Hirsch, the semantic function of a sub-sentential expression does not play an explanatory role in determining the character of a sentence – the character of the sentence is not 'built up' or determined by semantic facts at the sub-sentential level. Rather, the character of a sentence is explanatorily prior – semantic facts for sub-sentential components are only recoverable afterward. This much was noted by Hawthorne (see chapter five, section 7). The result of this is that semantic facts about sub-sentential components are, in a context of utterance, trivial and uninformative. The most salient example of this is facts about the references of singular terms. Hirsch is committed to a deflationary view of reference (see chapter two, section 3.1). On this view, there is no substantive, explanatory reference relation which links singular expressions to objects in the world. On a substantive theory of reference, the fundamental way in which language relates to the world is via the substantive reference relation. The fact that a whole sentence then relates to the world can be analysed as being a result of the way the various sub-sentential components of that sentence relate to the world in combination. Such a view is, as mentioned, rejected by a deflationary view of reference. On this view, if it is true in a context of utterance that '*a* exists' in

L , then it trivially follows in L that ‘the expression ‘ a ’ refers’ is also true. This is merely an additionally true sentence and gives us no more explanatory insight than originally knowing the truth of ‘ a exists’. Thus on this view, the fact that in a context of utterance, a sentence ‘ S ’ has the particular truth conditions it does is not in any way explained by semantic facts about the references of singular expressions in ‘ S ’. Rather, on Hirsch’s construal of language, *the way that language is related to the world is at the level of a whole sentence*. What I mean precisely by this will be explored in chapter 8, where I will conduct an argument against Hirsch based on this point.

However, for now I want to argue that on a proper understanding of Hirsch’s views on language, it would not be possible for Sider to intelligibly ground the notion of eligibility constraints on quantifiers, for these conditions are simply trivially fulfilled in each of the alternative languages. The important point to note is that they are trivially fulfilled solely in virtue of Hirsch’s views on *language* – it implies nothing about Hirsch being committed the substantive, metaphysical claim that the world lacks quantificational structure, as Sider seems to think it does. There are two ways in which Sider attempts to cash out the idea of eligibility constraints of quantifiers, and in each case it will be seen that the explanation presupposes that there is a relation between language and the world that is more fundamental than the whole-sentence, i.e., in terms of the ways that sub-sentential components are related to the world. Thus to argue with such assumptions in mind is to miss a key aspect of Hirsch’s views on language.

Let’s now consider the ways in which Sider tries to cash out an explanation of what the eligibility constraints for quantifier expressions are. In his 2001, as discussed in section 2.1 of this chapter, he suggested that we can ground the notion of a natural quantifier meaning by positing a domain of ready-made objects in the world which the quantifier ranged over. However, Hirsch simply argued that if it is true that ‘there exists a table’ is true in L , then it also trivially follows in L that ‘the domain of ready-made objects contains a table’ (See section 3 of this chapter).

There is a way of cashing out how the expression ‘there exists a table’ could be true even if the domain of ready-made objects didn’t contain a table, but this, I suggest, would require the notion of a substantive reference relation. Suppose that the nihilist picture of reality is right – on this view then there would no table in the domain of ready-made objects. However, let us also suppose that charity

determines that the CSO-sentence ‘there exists a table’ is true. Now, one way of explaining how this could be the case would be the following. Quite clearly, the word ‘table’ in the CSO-language functions *syntactically* as a singular expression, but in reality there is no one-one reference relation between the word ‘table’ and some single object in reality (since we presume, for this example, that the domain of ready-made objects contains only simples). What the nihilist could say then, is that whilst the word ‘table’ syntactically functions as a singular term in the CSO-language, it actually is a group referring term which refers to simples arranged table-wise. Thus we might propose that the CSO-language quantifier is such that when it is combined with singular expressions, those singular expressions actually function as group referring terms. In this sense we can say the CSO-language doesn’t carve reality at the joints because in a language that *did* carve at the joints syntactically singular terms should only refer to singular objects in reality.

However, none of this matters if we do not recognise a genuine substantive relation between words and objects. From the CSO-language perspective, if ‘there is a table’ is true, it follows trivially in the CSO-language that the following is also true: ‘the singular expression ‘table’ refers to an object’. This follows not from any assumptions about the world and what it contains but simply from whether or not ‘there is a table’ is true in the CSO-language. Thus if it is a condition for a sentence of the form ‘*a* exists’ to carve at the joints that the syntactically singular expression ‘*a*’ refers to a single object, then given the deflationary theory of reference, if ‘*a* exists’ is true then this condition is *trivially fulfilled* in the language in question – and thus the sentence can be truly said, in that language, to carve at the joints. Therefore, we can say that the notion of joint-carving, if it is to be cashed out in terms of there being an objective domain of ready-made objects, cannot settle the question of which language carves at the joints, because in each language, it is true to say, of that language, that it carves at the joints. Thus the debate over what language carves at the joints will depend on what language one is speaking, and hence will *itself be a verbal debate*.

We can further note that this can be shown to be the case irrespective of the assumption of a language-independent domain of ready-made objects. If, as Hirsch’s view on language supposes, there is no substantive theory of reference, then language doesn’t ‘hook’ onto reality at the level of words. Thus if one’s explanation for why one sentence better carves at the joints than another involves appeal to the way that the sub-sentential expressions of those sentences hook onto

objects in reality, then that explanation fails because there is no ‘hooking onto’ occurring. Furthermore, denying that there is a substantive theory of reference creates difficulty for the intelligibility of presupposing that there is a language independent domain of ready-made objects by which to assess whether certain sentences carve better at the joints than others, for any description of the elements in that domain must be done in a language. Thus if we presuppose that the domain of ready-made objects contains only simples, that will have to be true in the CSO-language. Since that is actually false in the CSO-language, the presupposition fails. This again is just to repeat the point that the question of what the domain of ready-made objects contains will have different answers in different languages.

The other way Sider tried to cash out the idea of quantifier naturalness, discussed in section 5 of this chapter, was with the notion that all sentences of the form ‘something is ϕ ’ express similar fact types where some *thing* is in *some way*. It was remarked that from the nihilist perspective, the CSO-sentence ‘there is a table’ actually expressed a fact where some *things* (simples) were arranged in some way (table-wise). But when the CSO says ‘there is a simple’ they *are* expressing a fact where something is in some way. Thus the CSO uses sentences of the *same form* to express facts of *different types*, and this is less natural and eligible than having all sentences of the form ‘something is ϕ ’ as expressing a fact about some *thing* being in some way.

However, from the perspective of the CSO-speaker, their sentence ‘there is a table’ *does* express the fact that there is a table – i.e. a fact that there is something in some way. This follows from the idea that if a sentence is true, then a statement of what fact it expresses is just a trivial disquotation of the original sentence: ‘the sentence ‘ ϕ ’ expresses the fact that ϕ ’. Since the character of ‘there is a table’ is such that the sentence is true in the CSO-language, then it trivially follows that ‘there is a table’ expresses the fact there is a table’ is also true in the CSO-language. It also follows that ‘there is a table’ expresses a fact where something is in some way’ is also true in the CSO-language. Thus from the CSO perspective, it will appear as if their quantifier has the most natural meaning. This will be the case with each language. This follows not from any assumptions about the structure of the world itself, but simply from assumptions about language and how it relates to the world. If Hirsch is correct in this view, then attempts to ground the eligibility constraints of sub-sentential expressions in terms of substantive relations of ‘joint-carving’ will not be successful, because all such conditions must

be stated in the languages in which they are supposed to apply, and those conditions will be trivially fulfilled in each language.

The result of this is that Sider's attempt to ground the notion of eligibility constraints, insofar as they seem to rely on an approach to language that Hirsch doesn't share, can seem to be question begging. If we assume Hirsch's approach to language is correct, Sider's arguments are not effective. Thus in order for Sider's argument to have any force we have to assume that Hirsch is mistaken in some key aspects of his approach language. Sider did attempt to show this by arguing that the meanings of some expression were constrained by naturalness, and he argued that this could be extended to quantifier expressions. However, Hirsch's competing argument that use can trump eligibility was also seen to be fairly convincing to the extent that Sider seems to accept that in some cases it could be true. Thus, each side, to some extent, seems to be entrenched not only in their meta-ontological views but also in the meta-semantic views which seem to support those meta-ontological views.

The question then is how to proceed with the debate. In the next chapter I will take up a claim I have made here – that Hirsch's views on language are such that the primary way in which language relates to the world is at the level of the whole sentence. In the next chapter I will put forward an argument to show that even if Hirsch is right about this, there can still be room ontological debate that isn't merely verbal. Thus Sider and Hirsch appear to be at an impasse, brought about, I think, by Sider failing to recognise some key aspects of Hirsch's views on language. In the next chapter then I offer a possible way out of this impasse by arguing with Hirsch on his own terms.

Chapter eight: A final problem for Hirsch

1: Introduction

In the previous chapter I argued that Hirsch's views on language means that for him, language relates primarily to the world at the level of the sentence. I then argued that Sider doesn't fully appreciate this point, and as a result his attempts to explain how one language will carve better at the joints than another didn't appear to be successful. Because for Hirsch language relates primarily to the world at the level of the sentence, the conditions that Sider tries to set down to show how a language might have the most natural meaning for its quantifier expressions are trivially fulfilled in each language. The result of this was that Sider's account ultimately failed to convince. In this chapter I am going to develop an argument which does pay attention to the fact that for Hirsch language relates primarily to the world at the level of the sentence. First of all I will explain a bit more about what I mean by saying that language relates primarily to the world at the level of the sentence, by contrasting it with a view whereby language is related first and foremost to the world at the level of individual expressions. The result of this will be to show that for Hirsch, in order for a sentence to be true by virtue of a particular fact, nothing about his views on languages then implies that there must be a kind of structural similarity between the sentence and the fact. This is in contrast to a more traditional correspondence theory of truth which does require a kind of structural similarity – i.e. a correspondence – between true sentence and fact. I will then make the claim that given that this is the case, we can still intelligibly suppose that the facts themselves might have a kind of language independent structure which might align with one or another of the first-order ontological positions that have been considered in this thesis. Thus Hirsch's views on language, I argue, leave room for intelligible speculation about ontology.

Against this criticism however, it will be seen that Hirsch has a ready-made reply – and that is to argue that outside of any language, the facts themselves are unstructured. On the surface this appears to be a metaphysical claim about the nature of facts, but Hirsch does not provide much more detail on what facts are supposed to be, other than that they are unstructured and thus cannot have the kind of objective structure necessary to support one ontological thesis rather than another. I will explore this claim in detail, and argue that if facts are unstructured, the notion of structure devolves down into a question about the kind of structure

the world itself has. I will argue that Hirsch's position, if it is to maintain a semblance of realism, cannot allow worlds to have the kind of structure needed to support objective *de re* modal similarity between worlds, and hence can only exhibit qualitative structure. I will argue then that Hirsch's position implies that the world, outside of language, is monistic. This is an incoherent place for Hirsch to end up in since monism is an ontological thesis that rivals the kinds of debates that Hirsch is claiming are merely verbal. Against this position Hirsch has two options – he can retreat into thoroughgoing Carnapianism by abandoning any semblance of realism or anti-verificationism, or he can adopt a third position which takes the world to be nothing more than a plurality of primitive unstructured facts. I will explore how this thesis works in conjunction with Hirsch's central claim about debates in physical object ontology, showing how this position leads Hirsch into an obscure position that is difficult to make sense of. Thus this chapter will demonstrate that Hirsch is faced with a choice between incoherence, a retreat into thoroughgoing Carnapianism, or obscurity. In addition to this, I then offer a final option for Hirsch – he can maintain his realism (i.e. not retreat into Carnapianism) and avoid problematic commitments if he adopts an epistemicist position. However, this again is a further addition to his position which takes it beyond mere semanticism, and I argue that he then faces additional explanatory burdens.

2: Language relates primarily to the world at the level of sentences

I have claimed that for Hirsch, language relates primarily to the world at the level of the sentence, but what exactly do I mean by this? First consider a contrasting position whereby language relates to the world first and foremost at the level of individual expressions via a substantive reference relation, and the ensuing correspondence theory of truth that is based on this primary relation. On this view, the reason that a sentence is true by virtue of a particular fact is explained by a correspondence relation between the sentence and the fact which is generated by the underlying referential relations between sub-sentential expressions in the sentence and objects in the world. The resulting picture is one where what is required for truth is that there is a structural similarity between the sentence and the fact that it is depicting. Thus on this view, a true sentence will wear its ontological commitments on its sleeve. If the sentence 'there is a table' is true, it will be true because there is a single object which has the property of being a table

existing in language independent reality⁷⁶. Thus on this view of language we shall say that it is *required* for a sentence to be true that it does indeed match the structure of the fact which it is depicting.

By contrast, this is not required for a sentence to be true on Hirsch's understanding of language. First of all, let us recall that Hirsch's position requires that the truth-value of a sentence is explanatorily prior to the references of the sub-sentential components of that sentence (chapter two, sections 4-4.3). Thus there is no bottom-up explanation explaining how a particular sentence gets to be true of a particular fact. Furthermore, on this approach the content of a sentence is not determined compositionally by the underlying meanings of the sub-sentential expressions. This is because on this view, the character of a sentence, which determines the content of a sentence in every context of utterance, is assigned at the level of the whole sentence. Prior to the assignment of character, a sentence is, on this view, an un-interpreted string of meaningless symbols. Thus the content of a true sentence about the external world on this view is determined only by its character, which is not constrained by the underlying referential relations of the sub-sentential expressions in that sentence. This is why I say that for Hirsch language relates primarily to the world at the level of the sentence. There is no required correspondence between the structure of the sentence and the fact which it is a depiction of because the structure of the sentence does not play a role in determining the content of a sentence.

I think an interesting implication of this position is that if it is not required that a true sentence's structure match the structure of the fact it is depicting, then the fact may or may not have a different kind of structure to the one implied by the structure of the sentence, or it may indeed have no language independent structure at all. On this view of language, all we can talk about is the structure of the true sentence, which would then appear to carve up the fact in thought and language in a way which might not actually match the structure of the fact, because a correspondence of structure is not required. The implication of this is that two true sentences with different structures could be true of the same fact. This is what Hirsch wants to be able to say because that allows him to say that debates in physical object ontology are merely verbal – the different positions do not actually differ with respect to the facts they depict; they only differ with respect to the

⁷⁶Different theories of course will treat predicates differently depending on one's favoured metaphysical understanding of properties.

required structure or form a sentence must take in each language to be a true sentence of a particular fact. Thus the nihilist sentence ‘there are simples arranged table-wise’ and the CSO-sentence ‘there is a table’ would, in virtue of having the same character, be true of the same fact in reality even though they have different structures. That Hirsch has something like this in mind is confirmed by the following quotation:

...the basic idea of quantifier variance can be nicely formulated by saying that the same (unstructured) facts can be expressed using the different concepts of “the existence of a thing”, that statements involving different kinds of quantifiers can be equally true by virtue of the same (unstructured) facts in the world (Hirsch, 2002: 59)⁷⁷.

Consider then that for Hirsch, the nihilist concept of ‘the existence of a thing’ is different from the CSO concept, so that the sentence ‘there are simples arranged table-wise’ can be true of the same fact that ‘there is a table’ is true of, owing to the fact that they are using different kinds of quantifiers (according to how the different ‘meanings’ of quantifiers in each language for Hirsch appear to be due to different rules for sentences involving quantifier expressions being operative in different languages – see chapter two, section 2.2).

Thus it seems to me that for Hirsch there isn’t a requirement for there to be a structural similarity between a true sentence and the fact of which it is a depiction. This is borne out first by his view that the content of a sentence is determined by the character of a sentence, and that there is no bottom up explanation for why a sentence is true in terms of the references of sub-sentential expressions in that sentence. Secondly it is borne out by the fact that Hirsch himself admits that two differently structured sentences can be true by virtue of the same fact in reality. I argued earlier that a correspondence theory of truth suggests that a true sentence about the external world wears its ontological commitment on its sleeve, as in order for it to be true, the structure of the fact which it is depicting has to match the structure of the sentence. However, since that is not required here, a sentence being true of a particular fact does not seem to suggest that the fact has a particular kind of structure, or indeed any structure. This, I argue, means that it is thus intelligible to suppose that the structure of the language independent fact

⁷⁷ Please note that I am not ignoring Hirsch’s use of the term ‘unstructured’ yet, but am reserving discussion of this point for the next section.

could actually match one of the true sentences but not the other, or it might have no language independent structure (as Hirsch seems to argue here – I will consider this in the next section). Thus we could still intelligibly speculate about the actual structure of the fact even if differently structured sentences are true of that fact, because the truth of those sentences does not require that the fact itself have any particular kind of structure. Thus I think it remains intelligible to suppose that one ontological position might be correct because it could match the actual structure of the fact, in which case there is still some scope for substantive ontological debate (though at this point it is not clear how we will frame such debates).

To summarise the discussion thus far then – it seems that if there are facts that stand in a truth-making relation to sentences, it would appear that a traditional correspondence theory of truth requires the truth-making relation to be explainable by and dependent on more basic relations between the constituents of facts and the sub-sentential components of sentences. However, on Hirsch's approach, though there may be a truth-making relation between fact and sentence, there is no substantive explanation at a lower level that determines or explains that truth-making relation. However, given Hirsch's commitment to realism, it seems reasonable at this stage to suppose that the fact itself could nonetheless have a kind of structure which could vindicate a particular ontological position. Thus there *might* be a kind of structural similarity between a fact and particular sentence which is what ontologists will then be interested in discovering. In that sense it is then entirely plausible to think that although two sentences are true of the same fact, one of them depicts the structure of the fact more correctly than the other.

3: *The nature of facts*

It will not escape the attention of the reader however that Hirsch has a readymade reply against this kind of criticism – that is the argument that the facts themselves are 'unstructured' (see the quotation above) – it is only language that imputes this kind of structure 'onto' the world. Thus the claim that a fact *could* vindicate a particular ontological position is immediately blocked because Hirsch has made the claim that the facts themselves are unstructured. We might wonder exactly what argument Hirsch makes to establish such a claim, but there isn't one to be found in his writing. Quite simply the thesis fits his overall aim and he thus adopts it with no argumentative support. Regardless of that we can still consider the nature

of these facts and how they are supposed to fit into Hirsch's overall system.

Let's first of all consider that Hirsch seems to abandon fact-talk, which was present in his 2002, for talk of propositions as sets of possible worlds in his 2009. It is not made clear whether he wishes this kind of approach supersede his early talk of unstructured facts, or whether he thinks the two sorts of positions can be accommodated. If we are to be conciliatory toward Hirsch, presenting his position in the best possible light before criticising it, it is quite clear that facts cannot be the sorts of things that have objects, properties and relations as constituents, since again that would imply that facts are to be individuated in terms of their constituents (and their ordering), and then you would have to let back in the kinds of entities that ontologists are arguing over in order to make sense of what facts are. Indeed, if Hirsch is arguing that facts are unstructured then they cannot be taken to have any constituents at all.

Against this then we can consider the notion that facts, for Hirsch, are unstructured abstract objects that are correlated with true propositions. For Hirsch, every proposition is a set of possible worlds, but of course not all sets of worlds are true propositions. Thus Hirsch might say that correlated with every true proposition is a fact, and thus two intensionally equivalent sentences are true in virtue of the same fact. Thus Hirsch can retain the notion that facts are unstructured by correlating (but not identifying) them with sets of possible worlds. The reason Hirsch would not want to identify facts with sets of possible worlds is that sets can be said to have a structure in the sense that they contain elements – specifically possible worlds in this case, and thus sets will be individuated by their elements.

Before I now move on to discuss the argument to monism, I want to point out a number of issues Hirsch will have with this approach. First of all, to claim that facts are unstructured abstract entities seems to be an ontological thesis, and so Hirsch isn't really showing via language alone that the debates in physical object ontology are merely verbal – he is relying on the truth of some ontological claim. It then remains to be seen in what language the truth of this ontological claim has to be stated in, since any premise that Hirsch requires presumably has to be state-able in some language. However, the claim cannot obviously be true in the CSO meta-language, for example, since the CSO takes the object-level sentences 'there

is a table' and 'there are simples arranged table-wise' as being made true by different structured facts, despite being intensionally equivalent. It is thus not clear how insisting that facts, outside of language, are unstructured abstract entities does not amount to bare refutation of the CSO position, rather than merely showing the CSO position to be nothing more than a linguistic decision to speak a certain way – since the CSO position requires that facts be structured entities, and Hirsch is insisting that they are not in order to establish his central claim. I have argued at length how this particular observation is problematic for Hirsch in chapter six so will not repeat it here, but it is interesting to note this particular issue arising for Hirsch once again here.

4: The argument to monism

Thus far I have argued that in order to accommodate Hirsch's claim that facts themselves are unstructured with his adoption of the thesis that propositions are sets of possible worlds, that facts must be abstract entities that are correlated with true propositions. This neatly removes the worry that the question of a fact's structure might cause a problem for Hirsch in establishing that the debates in physical object ontology.

However, I now want to argue that the issue of structure does not go away – it merely gets devolved down to entities that are now doing a lot of work for Hirsch – possible worlds. Why is this? Consider that a true proposition is a set of possible worlds. The possible worlds in that set will all be similar to each other in a relevant way. For example, worlds at which 'Donald Trump is president' will all be similar in a particular way (though they will then differ in other respects). Now we might ask – in what way are the worlds at which the sentence 'Donald Trump is president' is true similar? The simple answer will be that all those worlds will contain an individual that is recognisably Donald Trump, and at each of those worlds the individual will have the property of being president. Other worlds will be similar to the extent that they contain Donald Trump, but at those worlds he does not have the property of being president. Other worlds will further differ in that they do not contain Donald Trump at all. So worlds can be thought of as being more or less similar in terms of *de re* modality – i.e. worlds can be similar or different with respect to the objects they contain (and other things such as the properties those objects possess at different worlds, and other things such as the relative locations

of those objects in those worlds). We can also recognise ways in which worlds can differ qualitatively, but might remain otherwise *de re* similar to other worlds.

We have already encountered, in chapter five, Hawthorne's idea that worlds can be qualitatively identical but be *de re* distinct. Although I rejected this sort of reasoning then as a way of criticising Hirsch, since Hawthorne's examples relied on the assumption of haecceities, we can nonetheless make sense of worlds which are *de re* identical despite minor qualitative differences. Consider for example a table, which we will call *Table*. When we make a *de re* modal claim about *Table*, the idea is that we are making a claim about the modal properties of the object itself, regardless of how it is referred to or talked about in language. Consider then that if I say '*Table* could lose one of its constituent simples', then if this is true, there is a possible world, *w1*, at which *Table* exists, but is composed of one less constituent atom. Assume the two worlds are identical in every other respect; it then seems to make sense to say that our world and *w1* are *de re* identical with respect to macro-physical composite objects, although there will be the slightest qualitative variation, since *Table* in *w1* will have a slightly lower mass/size. It also will be true that whilst this world and *w1* will be *de re* identical at the level of macro-physical composite objects, there will be a *de re* difference at the level of simples.

Consider now that the CSO-speaker will claim that this world and *w1* are *de re* identical (at the level of composite objects), whereas the nihilist will not recognise this world and *w1* as being *de re* identical, since the only kind of modal properties they recognise are those belonging to simples. Thus the nihilist will deny that this world and *w1* are *de re* modally similar on any level, whereas the CSO will recognise *de re* modal identities at a certain level of composition. However, there can't be *anything about the world itself which settles this issue*, since the dispute between the nihilist and the CSO is supposed to be merely verbal.

What does this mean for Hirsch? Well, just as facts could not be thought of as having language-independent structure that might support one ontological position over another, *worlds* cannot have the kind of structure needed to support genuine, language-independent *de re* modal similarity/dissimilarity between worlds that would vindicate one ontological position over another. This is because if the world did contain the structure needed to support genuine *de re* modal similarity, then it wouldn't be merely a matter of linguistic decision which worlds were more or less *de re* similar, which means it couldn't be a merely a matter of linguistic decision

what objects worlds contain. Thus just as Hirsch must argue that outside of language, facts are unstructured entities, he must *also* claim that outside of language the world cannot have the kind of structure needed to support genuine *de re* modal similarities between worlds.

But if the world cannot contain the kind of structure needed to support *de re* similarities between worlds, the only objective similarities between worlds can be qualitative similarities – *de re* similarity between worlds, somewhat paradoxically so, must be language-relative only.

Thus far then we have, by considering the idea of worlds being more or less *de re* similar, seen that Hirsch cannot allow worlds to have the kind of structure needed to support such language independent similarity. Thus, outside of language Hirsch cannot *allow for the notion that the world divides into objects* – since *de re* modal claims involve the idea of objects and their existence in different possible worlds. However, since Hirsch makes an explicit commitment to realism, he will have to admit of there being *some way* the world is independent of language, and thus the world's existence itself, as a language independent entity must be affirmed, but it cannot be such that it divides into objects. Presumably as a realist, he will also have to admit of there being some objective way in which worlds are more or less similar. Thus the only genuine similarity between this world and other worlds can be qualitative similarity. This appears to commit Hirsch to existence monism – the only thing we can positively affirm as existing is the world itself, as a single entity that exhibits qualitative variation.

The problem with ending up in this position for Hirsch is that monism itself is an ontological thesis that is in opposition to the ontological theses that Hirsch claims are a result of nothing more than linguistic decision. Thus Hirsch's position ends up as incoherent, since he appears to be tacitly committed to monism, whilst arguing that the dispute in physical object ontology are merely verbal. However, if Hirsch is committed to monism then the other positions in physical object ontology are mistaken in a substantive way – thus the issue is *not* merely verbal.

It will be noted that Hirsch's position runs into incoherency only because of his insisted commitment to realism. For example, the argument above has frequently made reference to the requirement for the world to be some way independent of

language – particularly the notion that the world could be objectively more or less similar to other possible worlds. One way of resisting such talk then is to insist that the notion of the world being some way outside of language is nonsensical. But to insist on this is to retreat from realism, and is thus to give up on one of the key commitments Hirsch makes to distinguish his position from a thoroughgoing Carnapianism. The threat of collapse into incoherency as shown here further demonstrates that Hirsch's modified light version of Carnapianism is untenable.

5: A way of avoiding monism – Tractarian factualism

In the previous section I discussed how Hirsch's modified light version of Carnapianism leads him to existence monism, thus making his central claim that the debates in physical object ontology are merely verbal incoherent. Against this view it could be argued that I have ignored a further position available to Hirsch which allows him to maintain a commitment to realism but which also avoids the claim that his position ends up committing him to monism.

It involves returning to the notion of unstructured facts, and offers a different account of the nature of these facts that I earlier talked about. The idea then is that Hirsch can adopt the Tractarian inspired view that the world itself consists of a plurality of primitive unstructured fact. This position allows Hirsch to avoid monism since now technically the world cannot be said to be a single, unstructured entity, since the world is a totality of particulars, albeit these particulars are *not* objects, but are instead primitive, unstructured facts. I will first explain how this position is supposed to work. I will then discuss two worries. The first is that this option once again does not seem to address the issues I raised in chapter six. Secondly, trying to fit together the notion of what I will refer to as *Tractarian factualism*, or *factualism* for short with the claim that the debates in physical object ontology leads to an obscure position that is difficult to make sense of. I will also then argue that whilst factualism in conjunction with Hirsch's merely verbal claim is not technically monism (given the fact that the world is taken to consist fundamentally of a plurality of particulars), it nonetheless is tantamount to monism, but with the added detriment of being explanatorily more obscure. Thus whilst adopting factualism might allow Hirsch to avoid a technical commitment to monism, he ends up in a position that is no more attractive given his stated aims.

Firstly, how are then to understand the notion of *factualism*? Such an idea has its historical precedent in Wittgenstein's *Tractatus*, and is summarised by Skyrms as follows:

We may conceive of the world not as a world of individuals or as a world of properties and relations, but as a world of facts – with individuals and relations being equally abstractions from the facts. John would be an abstraction from all facts-about-John; Red an abstraction from being-red-facts; etc. (Skyrms, 1981: 199).

We should first understand that for Hirsch, it cannot be a matter of merely 'conceiving' of the world as a world of facts as opposed to individuals – rather he must be *committed* to this kind of factualist metaphysics outright. This is because he is claiming that the question of what objects the world contains is merely a matter of linguistic decision. But if he is only doing that from a position where he can 'conceive' of the world as consisting of facts, then his merely-verbal claim is itself dependent on what conception of reality we choose to adopt (since I can equally conceive of the world as a world of individuals). If I decide to conceive of the world as individuals, then Hirsch's merely-verbal claim will not be effective since I do not share his factualist conception of the world. Thus Hirsch's merely-verbal claim would itself be merely verbal, which looks incoherent. It seems to me that his merely verbal claim cannot simply come out true according to the factualist conception, which itself is understood as merely another way we can conceive of the world – at least he would not want it to be this way.

Secondly we need to understand Skyrms claim that individuals are *abstractions* from facts. It is not clear in what Skyrms has said whether objects (I shall revert to the term objects rather than individuals), understood as abstractions from facts, are based on *similarities* between facts. The question then is what is the basis of this abstraction? Let's assume for now that abstractions are based on similarities between facts – the second question now is whether these similarities are objective or merely linguistically relative. Consider that in further characterising factualism Skyrms says that 'The way we break things up depends on what objects and relations we take as being most *generally* useful in characterizing the world'

(Skrirms, 1981: 201). This is an ambiguous claim since it is compatible with us understanding there being objective similarities between facts, but that we tend to focus on the ones that that are most useful for is, or that fact similarity is dependent on our interest, and so fact similarity wouldn't be objective. The latter reading is compatible with Hirsch's view, since Hirsch can argue that objects are merely abstractions from facts based on our interests, and that there is no objective notion of fact similarity – i.e. there is no objective notion of what objects there are.

Let me explain this latter point in more detail. Essentially, I am claiming that if Hirsch is to adopt a factualist position, then he must make the claim that objects, as abstractions from facts, can only be based on subjective similarities between facts, not objective similarities. Why couldn't there be objective similarities between facts? Hirsch could not allow there to be objective similarities between facts since then the objective similarities would give rise to objectively existing objects, albeit these objects would then be understood as ontologically derivative entities, being nothing more than objective similarities between facts. However, even though the objects are now ontologically derivative, it would still be a substantive issue whether such objects existed, since it is a substantive matter whether there are such objective similarities between facts. Thus Hirsch must insist on a kind of factualism where fact similarity is taken to be purely subjective⁷⁸ – there can be no room for the kind of objective similarity between facts that would give rise to objectively existing objects as derivative entities.

I now want to consider the first problem with such an account. Even with this factualism in place, Hirsch will still need to argue that the sentences 'there is a table' and 'there are simples arranged-table-wise' are made true by virtue of the same fact in the world. Given that facts themselves are objective and primitive entities, here envisioned to play the role of truth-makers for sentences, then it would seem to be an objective matter whether two sentences were made true by virtue of the same fact or by distinct facts. Hirsch requires it to be the case that the sentences 'there is a table' and 'there are simples arranged table-wise' are made true by the same fact, but again this doesn't seem like it shows up the CSO

⁷⁸ Perhaps purely subjective here is wrong – since there would be an intersubjective agreement on fact-similarity. We might consider, for example, the intersubjective agreement of the Melville community (discussed in chapter 7) who claim that whales are fish.

position as a merely verbal positions – rather it looks like the CSO position cannot be right in any sense. I have already covered this issue in chapter six, but it does not appear to go away even when we consider that Hirsch could adopt a factualist position. Regardless, this is not the main issue facing Hirsch's factualist position.

Remember we are considering Hirsch's factualism as an alternative to monism and hence incoherency and a retreat in anti-realism and a thoroughgoing Carnapianism. It would thus seem that Hirsch's factualism would need to retain a semblance of realism, as opposed to being a dressed-up version of a thoroughgoing Carnapianism.

What must be objective and language independent according the factualist who wishes to maintain a semblance of realism? Consider that for Hirsch, we cannot have any language independent objective similarity between facts that would give rise to objects as derivative entities. However, as a realist, Hirsch will presumably want to maintain that the world exhibits nomic/causal features that are nonetheless independent of our thought and language⁷⁹. Given that, what could be the ground for such nomic/causal features on this approach other than objective similarities between the primitive unstructured facts?

However, the problem with this is that we track causal relations in the world in terms of the histories of objects. But how that is possible given that the objects we track, understood as fact similarities, are subjective and not objective is obscure and difficult to make sense of. This is because tracking the histories of subjective objects would infect the causal/nomic features of which we speak, leaving the objective causal/nomic features out of our cognitive reach.

Against this, Hirsch could argue that science can discover objective fact similarities and that is what explains the causal/nomic features of the world, but these sorts of objective similarities are not the same sorts of similarities that would give rise to objectively existing but derivative objects. But then this position is claiming that there is enough objective fact similarity to explain the causal/nomic features of the world, but that there is not enough objective fact similarity to generate objects. So this position ultimately claims that whilst there is an objective world, which is the totality of facts, such that there is enough fact similarity to give us a world with

⁷⁹ Whilst of course maintaining that such causal/nomic features are discoverable by us.

causal/nomic structure, the same world of facts does not contain any kinds of objective similarities to give us objects. Whilst this position is not technically monism, since the world is a plurality of particulars, the nature of these particulars is now made entirely obscure. The monistic position claims something similar – the world contains or exhibits causal/nomic features but no objects. The difference between the positions of course is that the factualist position takes as primitive a plurality of facts and the monist takes as primitive a single concrete entity. Beyond that both positions do not allow for existence of objective objects (beyond the world itself in the case of monism). Thus whilst Hirsch opting for a factualist approach does not entail a commitment to monism, it does involve opting for a position which has similar implications, with the added cost of being more obscure.

As such, Hirsch appears to be faced with a choice between a commitment to monism and thus incoherence, factualism and thus obscurity, or finally a retreat into thoroughgoing Carnapianism, showing that his modified light version of Carnapianism is untenable.

6: A final option – epistemicism

In the previous section I construed Hirsch as facing three options – he must either accept that his commitment to realism leads him either to an incoherent monism, or it leads him to an obscure factualism, or finally to give up on realism and retreat into thoroughgoing Carnapianism. However, a final option open to Hirsch might be to embrace a kind of epistemicism. He might say, for example, that although the world might contain genuine *de re* structure, or although the world might contain objective fact similarities, these sorts of things are out of cognitive reach and that there is no point in speculating on them. In this way, Hirsch avoids an outright commitment to monism or factualism, and can still retain the idea that there is nothing at stake between describing the world in a nihilist language or a CSO language – since as far as we can ever tell both languages are equally good at capturing what we can capture of reality – so there is nothing to decide between the languages other than linguistic decision. Whilst Hirsch could adopt this position, it is obviously not the position he originally intended to articulate since now it is relying on an epistemological claim about the limits of what we can know which was not present in his original formulation. Thus he would no longer be defending

a merely-verbal claim – it would be a merely verbal claim based on an epistemological claim.

Against this position however, the ontologist can simply take issue with Hirsch's bleak epistemicism and argue that whilst we can't really step outside language to check if our ontological theories are true once and for all, we could still argue over what theories are more or less likely to be correct. The way to do this is to do what ontologists have hitherto already been doing - adopt an epistemological approach which argues that theory selection based on theoretical or explanatory virtues is truth-tracking. As L.A. Paul explains:

The theoretical desiderata we use to choose a theory include simplicity, explanatory power, fertility, elegance, etc., and are guides to overall explanatory power and support inference to the truth of theory. A scientific realist should take such desiderata to be truth-conducive, since it is hard to see how such desiderata can lead us to truth if they are merely or even mainly pragmatic virtues. If such theoretical desiderata are truth conducive in science, they are also truth conducive in metaphysics (and in mathematics, and in other areas). The main point I want to make here is that if the method can lead us to closer to the truth in science, it can lead us closer to the truth in metaphysics (Paul, 2012: 21).

The idea then is that the modern metaphysician will approach metaphysical/ontological problems with a similar epistemological method to the scientific realist. This of course is by no means a conclusive way of establishing that ontological debate is substantive. However, if Hirsch wishes to avoid a retreat into Carnapianism, and wants to avoid further problematic commitments to monism or factualism, then he can only take the epistemicist option. As well as being an addition to the merely-verbal claim, Hirsch would also need to offer an argument against the notion that theoretical virtues in metaphysics are not truth-tracking whilst they are truth-tracking in scientific theories in a way that chimes with his commitment to realism.

7: Summary of the argument in this chapter

In this chapter I argued that Hirsch's views on language mean that a fact can stand in a truth-making relation to a sentence without there being any further underlying

explanation of how that truth-making relation is established via more basic relations between the constituents of facts and sub-sentential components. This meant that a sentence could be made true by some fact even if it didn't match the structure of the fact.

I then argued that this means there is still room for intelligible speculation on what the structure of the facts are like, and they could be structured in a way that vindicated a particular ontological position. Against this however I pointed out that Hirsch specifically commits to facts being unstructured entities.

I then argued that for Hirsch, the best understanding of the nature of facts is that they are abstract unstructured entities that were correlated with true propositions, understood as sets of possible worlds. I then made an argument suggesting that Hirsch would have to abandon the notion that the world could have the structure needed to support genuine *de re* modality since he would need *de re* modal claims to be true only relative to some language. This, I argued, meant that Hirsch was committed to monism, which is incoherent given that he is trying to establish that debates in physical object ontology are merely verbal.

Against this I argued he could give up his claim to realism and retreat into thoroughgoing Carnapianism, further showing his modified light version of Carnapianism to be untenable. However, I then considered a further option – Hirsch could adopt a factualist picture of the world, with the further commitment that fact similarities could only be understood as relative to some language. This commitment became a problem when we considered Hirsch's commitment to realism and how he could account for the objective causal/nomic structure in the world without having objective fact similarities. I thus argued that the factualist position, in conjunction with Hirsch's merely verbal claim, and a commitment to realism, ultimately led to an obscure position.

Finally I argued that Hirsch could adopt an epistemicist approach, but argued that this again shows that his original position as presented is untenable as it would need the further addition of an epistemological claim. Against this, I then argued that Hirsch was faced with a further explanatory burden of showing why theoretical virtues were not truth-tracking in metaphysical theories but were truth-tracking in scientific theories.

8: General conclusion

The purpose of this thesis has been to critically assess the effectiveness of Eli Hirsch's semanticist deflation of the debates in physical object ontology. Ultimately, the thesis has demonstrated that Hirsch presents a Carnapian-inspired position which is committed to the core claim that the disputes in physical object ontology are merely verbal. But I have argued that Hirsch intends to put forward a modified light version of Carnapianism which seeks to avoid the anti-realist and verificationist implications of Carnap's views. Thus, I understand Hirsch as trying to situate himself as occupying a middle ground between substantive approaches to ontological issues and a thoroughgoing anti-realism/verificationism. Thus Hirsch seeks to vindicate common-sense, and seeks to be committed to a 'robust' realism. Further to this he also construes his position as anti-verificationist, and as applying only to the debates in physical object ontology (thus leaving room for some substantive ontological debates such as the dispute between nominalists and Platonists). The attraction of Hirsch's position, as far as I see it, is that it offers those who are sceptical of the debates in physical object ontology a way of deflating those debates, whilst still avoiding what might be seen as unattractive commitments to anti-realism, verificationism and detachment from common-sense. Unfortunately for those people, this thesis has uncovered considerable tensions within Hirsch's position, showing that his core claim that the debates in physical object ontology often runs into problems which can only be resolved by retreating toward a thoroughgoing Carnapian position. What this demonstrates is that Hirsch's modified light version of Carnapianism is not sustainable – either the debates in physical object ontology are substantive or something like a thoroughgoing Carnapianism is correct. There is no room for Hirsch's middle ground. In what remains of this section I will briefly remind the reader of the contents of the previous chapters before offering a final summary.

In the second chapter I offered a thorough explanation of Hirsch's position by analysing closely the interrelated claims of three of his papers. The result of this chapter, I think, throws up interesting conclusions that others haven't fully appreciated or even noticed. One of the main ones is that Hirsch, amongst other commitments in the philosophy of language, is committed to a deflationary theory of reference. In the chapter as well I offered an analysis of what quantifier variance amounted to and argued that this ultimately amounted to the claim that in different languages there can be different rules governing sentences involving

quantifier expressions, with the result that the existential claims have different truth-conditions in different languages. Thus the disagreement over whether a particular existential claim is true can be shown to be merely verbal if that claim has different truth conditions in different languages owing to nothing other than the different semantic rules operative in those languages.

In this chapter I also tied up the implications of Hirsch's views and arrived at the argument that for Hirsch, the character of a sentence, which determines its content in a context of utterance, is ultimately determined by how that particular sentence is used in a particular linguistic community. Thus for Hirsch, it is how a whole sentence is used in a linguistic community that determines its truth-conditions. I also argued that before a sentence has been assigned a character it is a meaningless string of symbols, and thus another explanation for how the debates in physical object ontology can be merely verbal is by showing that different (meaningless) sentences can be used by different communities in all the same contexts (such as 'there is a table' and 'there are simples arranged table-wise'). Thus differently structured sentences would then have the same character across the different languages, and since the character of the sentence determines its content in every context of utterance, the sentences themselves did not differ with respect to the facts they expressed. Thus on Hirsch's view, intensionally equivalent sentences have the same content, even if they are 'structured differently'. In this chapter I also set out how Hirsch intended his position to be distinct from a thoroughgoing Carnapianism by his explicit commitment to the following views 1) he is committed to robust realism, 2) he is explicitly anti-verificationist, 3) he envisions his criticism as applying locally only to physical object ontology and finally 4) he seeks to vindicate or defend common-sense. In chapters 3-8 I then developed several lines of argument against Hirsch's main position, showing that Hirsch cannot maintain his central claim without giving up on his other explicit commitments, thus showing the position he presents as being untenable.

In chapter three I explored an argument put forth by Bennett who argued that if the debate between the nihilist and the CSO was merely verbal, this could only be the case if the sentence 'LP*' ⁸⁰ was necessarily true in the CSO-language in virtue of being analytic in that language. In that chapter I expanded upon some blind spots in Bennett's argument. First of all I offered an explanation (which I think needs to be given) to show why, if the dispute is to be verbal, that 'LP*' couldn't be

⁸⁰ 'If there are simples arranged table-wise then there is a (numerically distinct) table'

analytic. This much appeared to be assumed by Bennett but it is a crucial horn of the dilemma. If it wasn't analytic it would have implied that the CSO position was effectively correct because the only way 'LP*' could be necessary, aside from being analytic, was if it described a necessary connection between distinct objects in the world. Against this then, Bennett argued that 'LP' (if there are simples arranged table-wise then there is a table) couldn't be analytic because it ignores the numerical distinctness claim that the CSO speaker makes – the CSO speaker does not regard the expression 'table' as referring to the same object(s) as 'the simples arranged table-wise'. She then argued that we couldn't modify 'LP' to 'LP*' because that would imply we could define things into existence.

Against this I argued that Bennett here hadn't paid sufficient attention to Hirsch's views on reference. Drawing on my discussion in chapter two then, I showed that Hirsch could coherently maintain that 'LP*' was analytic in the CSO-language without thereby admitting that we can define things into existence. However, luckily Bennett offers another reason for thinking that 'LP*' can't be analytic in the CSO-language – the CSO speaker simply does not regard 'LP*' as analytic. I argued that this in fact was a more effective criticism of Hirsch's position. Hirsch's view is overly focused on establishing that the nihilist language and the CSO-language (for example) are intensionally equivalent – but here he only focuses on object-level claims. This, I pointed out, ignores the ways in which ontological positions also have a stake in certain meta-level semantic propositions. Thus Hirsch cannot simply assume what the truth-values of these are without begging the question against at least one side in the debate. This is because he is foisting a certain meta-level claim upon a speaker of a certain language when that Hirsch ends up analysing a debate that no-one is actually having. Against this, I argued that Hirsch could argue that there was no difference between object-level necessary truths that were analytic or synthetic. This, I argued, threatened Hirsch's commitment to anti-verificationism. A further argument I explored on Hirsch's behalf was to argue for no clear-cut distinction between analytic and synthetic truths. This, I argued, was a step too far for Hirsch as it means there is no clear-cut distinction between merely verbal debates and substantive ones, and Hirsch is firmly committed to some debates being definitely substantive and others being definitely merely verbal. Thus the analytic problem posed a serious problem for Hirsch's position.

In chapter 4 I offered an original solution to the analytic problem which I think Hirsch could make, but showed that in many ways the cure would be worse than the disease – so Hirsch would either be stuck with the analytic problem or the equally problematic solution to the analytic problem. Let me remind you of some of the key points. I argued firstly that Hirsch' could simply argue that he isn't making an explicit commitment to whether 'LP*' needs to be analytic in the CSO language – all he needs to establish is that the nihilist can usefully interpret it as being analytic. Drawing in discussion from the previous chapter I argued that Hirsch could not take the route. I then argued that Hirsch could simply envision each side as speaking not only a different object-language, but also a different meta-language. Thus the sentence 'the sentence 'LP*' is analytic' would be true or false depending on which meta-language one was speaking. This shows that the debate over the analyticity of 'LP*' is itself merely a verbal dispute. I then explored the implications of this position, arguing that Hirsch, whilst allowing each side in the debate to maintain an illusory distinction between synthetic/analytic necessary truths, would have to be committed to the claim that all necessary truth are in fact analytic. This was shown to a problem for Hirsch because it meant he was committed now to anti-realism, thus showing that this particular solution to the analytic problem leads Hirsch to giving up on one of his key commitments, and represents a retreat to a thoroughgoing Carnapianism. In addition, I argued that giving up on realism meant that Hirsch was also giving up on any serious defence of common-sense, and that he would also have to give up on the claim that his merely-verbal claim applied only to debates in physical object ontology. Thus by the end of chapter four it is beginning to look like Hirsch's modified light version of Carnapianism is untenable.

In chapter five I then moved on to discuss several objections to Hirsch's position raised by Hawthorne. These divided into what Hawthorne called 'intensional issues' and 'hyperintensional issues'. In this chapter I mainly showed that the intensional issues raised by Hawthorne could easily be met and so didn't pose an issue for Hirsch. I then moved on to discuss the hyperintensional issues. One interesting point Hawthorne raised was that the success of Hirsch's project rests upon scepticism toward bottom-up semantics, and that if a bottom-up approach could be seen to be explanatorily illuminating then it would create a problem for Hirsch's verbal claim. I argued that this was a good point, but it didn't really demonstrate a decisive refutation of Hirsch's position. Finally, I discussed a point that Hawthorne

raises but didn't develop into an actual argument against Hirsch. This point was that metaphysicians often make use of hyperintensional operators in order to track structural features of reality. However, Hawthorne points out that for Hirsch, hyperintensional operators can play no role in determining the content of a sentence, owing to his 'intension-centric' approach to semantics. Whilst this was an interesting point to bring up, Hawthorne failed to demonstrate why this constitutes a problem for Hirsch and not for the metaphysician.

This then led me onto to chapter six, where I took up the point raised by Hawthorne and developed it into proper argument showing why hyperintensionality poses a problem for Hirsch. I called this problem the H-problem. The H-problem can be summarise as follows: Hirsch himself recognises two levels of content – coarse grained content that is captured at the level of the intension of the sentence and fine-grained intentional content. He owes us an account of this distinction in content. I argued that this distinction in content could be seen to be hyperintensional since intensionally equivalent sentences which purportedly differed in their fine-grained content could be prised apart using hyperintensional operators. I then argued that Hirsch owed us an account of this distinction in content. Essentially there could only be two accounts of hyperintensional content – it either arises from our different representations of the same facts in reality, or it arises because we are tracking structural features of reality that cannot be gleaned at the level of intensions. I argue this presents a problem for Hirsch because whilst Hirsch will need to claim that hyperintensional distinctions in content are merely representational in order to establish that the dispute between the CSO and the nihilist is merely verbal, this again further ignores the ways in which the ontologist has a stake in whether hyperintentional distinctions in content are representational or due to the world itself. Specifically, the CSO speaker will argue that some of their object-level sentences, though intensionally equivalent, are true by virtue of distinct facts. Hirsch cannot thus assume they are not true by virtue of distinct facts without begging the question. This then was the H-problem, and I noted it's similarity to the analytic problem. I then explored potential solutions – the most promising seemed to be that Hirsch could simply interpret the dispute over hyperintensionality as a further verbal dispute (like the proposed solution in chapter four). This, however, leads us back into a commitment to Tractarian analyticity and all the problems that brings for Hirsch's position. Thus by the end of chapter six we have seen that there are two distinct problems that emerge for

Hirsch's view which push him toward a thoroughgoing Carnapianism as the only way forward, thus showing that his modified light version of Carnapianism is not tenable.

In chapter seven I then moved on to discuss the debate between Sider and Hirsch. In this chapter I argued that the debate between Sider and Hirsch had two distinct phases. The first phase concerned a dispute over meta-semantics. Sider argued that in order for Hirsch's position to work, there would have to be no external constraints on the meanings of quantifier expressions in the various languages that Hirsch envisions – the meanings would be determined by use alone. Against this, Sider argued that there are external constraints on quantifier expressions and thus the possible languages that Hirsch envisions couldn't exist. Against this Hirsch argued that use could trump eligibility.

The second phase of the debate could be marked by Sider's concession of this point. However, Sider then argues that even if use does sometimes trump eligibility, we could speak a language where it was stipulated that we would be using quantifiers with their most natural meaning. Against this Hirsch argued that the debate over which language that was would be trivially true. I argued in favour of Hirsch on this point – showing Sider's reasoning to rest on question begging and a misunderstanding of Hirsch's views on language. Like Bennett, Sider does not pay sufficient attention to Hirsch' top-down approach to language, which I argued meant that for Hirsch, language relates primarily to the world at the level of sentences, and not at the level of sub-sentential expressions.

This led me to my final argument against Hirsch as presented in this final chapter. His commitment to language relating to the world at the level of the sentence meant that we could enquire as the structure of the facts in a way that might vindicate a particular ontological position. Against this, Hirsch argued that facts in themselves were unstructured. I explored the implications of this claim, showing that Hirsch, if he wishes to maintain his realism, ends up committed to monism or an obscure form of factualism. Against these two options, Hirsch could either retreat into thoroughgoing Carnapianism or adopt an epistemicist approach, both of which signify a retreat from his original position.

All things considered, this thesis demonstrates some problematic tensions in Hirsch's position. He wants to claim that a certain subset of ontological debates are merely verbal, but he wants to do it in a way that avoids committing himself to

anti-realism or verificationism, and he always seems to completely avoid taking a stand on the issue of analyticity. In this thesis, it has been shown that Hirsch cannot avoid this issue, and that by paying attention to it, some serious problems for Hirsch's position have been uncovered.

Where does this leave physical object ontology and semanticism? Whilst this thesis has shown that Hirsch's light version of Carnapianism is untenable, it hasn't offered an assessment on the tenability of thoroughgoing Carnapian position. Thus whilst there still might be hope for a semanticist deflation of physical object ontology, it appears that those attracted to such a notion will have to look to adopt more faithful versions of Carnapianism, and embrace the anti-realist and verificationist approaches to deflating physical object ontology.

There are also other threats to the possibility of physical object ontology which also warrant some attention – the threat of epistemicism, the idea that we will never be in a position to gain actual knowledge about matters of physical object ontology, still looms large. Additionally, there are calls for metaphysics to be more scientifically respectable – for example Ladyman and Ross in their 2007 argue that the kind of picture of reality that ontologists work with is outmoded with regard to current science, and thus any conclusions reached in ontology would be without merit. Thus those who take physical object ontology seriously still have work to do to secure the legitimacy of their enquiry.

It should be remembered that there have always been sceptics regarding the possibility and intelligibility of metaphysical knowledge and the legitimacy of metaphysics as a respectable discipline. Kant of course famously argued that genuine knowledge of the world as it is in itself is impossible, and the logical positivists attempted to show that many purportedly meaningful metaphysical claims were meaningless. Metaphysics has always seemed to find a way to bounce back, only for new detractors to arise. The assault on physical object ontology then, from several fronts, can be seen as the latest incarnation of this battle. This thesis can be regarded as ammunition in that ongoing battle.

Bibliography

- Austin, J.L. "Truth." *Aristotelian Society supplement* 24, no. 1 (1950).
- Bennett, Karen. "Composition, Colocation, and Metaontology." Edited by David Chalmers, David Manley and Ryan Wasserman. Oxford: Clarendon Press, 2009.
- Blackburn, Simon, and Keith Simmons. *Truth*. Oxford: OUP, 1999.
- Bliss, Ricki, and Kelly Trogdon. "Metaphysical Grounding." *Stanford Encyclopedia of Philosophy* (2015 Fall Edition), 2015.
- Butterfield, Jeremy. "The Rotating Discs Argument Defeated." *British Journal for the Philosophy of Science* 57 (2006).
- Carnap, Rudolph. "Empiricism, Semantics and Ontology." *Revue Internationale de Philosophie*, 1950.
- Chalmers, David, David Manley, and Ryan Wasserman, . *Metametaphysics: New essays on the foundations of ontology*. Oxford: Clarendon Press, 2012.
- Davidson, Donald. *Inquiries into Truth and Interpretation*. Oxford: OUP, 1984.
- Effingham, Nikk. *An Introduction To Ontology*. Cambridge: Polity, 2013.
- Eklund, Matti. "Neo-Fregean Ontology." *Philosophical perspectives* 20, *Metaphysics* (2006).
- Goodman, Nelson. *Fact, Fiction and Forecast*. Cambridge (US): Harvard University Press, 1950.
- Hawthorne, John. "Superficialism in Ontology." In *Metametaphysics*, edited by David Manley, Ryan Wasserman David Chalmers. Oxford: Clarendon Press, 2009.
- Hirsch, Eli. "Language, Ontology, and Structure." *Nous* 42, no. 3 (2008).
- Hirsch, Eli. "Ontology and Alternative Languages." In *Metametaphysics*, edited by David Chalmers, David Manley and Ryan Wasserman. Oxford: Clarendon Press, 2009.
- Hirsch, Eli. "Physical-Object Ontology, Verbal Disputes, and Common Sense." *Philosophy and Phenomenological Research* LXX, no. 1 (2005).
- Hirsch, Eli. "Quantifier Variance and Realism." *Philosophical Issues* 12 (2002).
- Jespersen, Bjorn. "How hyper are hyper propositions?" *Language and Linguistics Compass* 4, no. 2 (2010).

- Kripke, Saul. *Wittgenstein on Rules and Private Language*. Cambridge (US): Harvard University Press, 1982.
- Ladyman, James, and Don Ross. *Every Thing Must Go: Metaphysics Naturalized*. Oxford: OUP, 2007.
- Lewis, David. "Putnam's Paradox." *Australasian Journal of Philosophy* 62, no. 3 (1984).
- Lowe, E.J. "In defense of medium-sized specimens of dry goods." *Philosophy and Phenomenological Research* 67, no. 3 (2003).
- Lowe, E.J. *The Four-Category Ontology: A Metaphysical Foundation for Natural Science*. Oxford: Clarendon Press, 2006.
- Lowe, E.J. "The Problems of Intrinsic Change: Rejoinder to Lewis." *Analysis* 48, no. 2 (1988).
- Markosian, Ned. "Brutal Composition." *Philosophical Studies* 92 (1998).
- Merricks, Trenton. *Objects and Persons*. Oxford: Clarendon Press, 2001.
- Ney, Alyssa. *Metaphysics: An Introduction*. Abingdon: Routledge, 2014.
- Nolan, Daniel. "Hyperintensional Metaphysics." *Philosophical Studies* 171, no. 1 (2014).
- Paul, L.A. "Metaphysics as modeling: the handmaiden's tale." *Philosophical Studies* 160 (2012).
- Price, Huw. "Metaphysics after Carnap: The Ghost Who Walks?" In *Metametaphysics*, edited by David Chalmers, David Manley and Ryan Wasserman. Oxford: Clarendon Press, 2009.
- Sider, Theodore. "Criteria of personal identity and the limits of conceptual analysis." *Philosophical Perspectives* 15 (2001b).
- Sider, Theodore. *Four-Dimensionalism: An ontology of persistence and time*. Oxford: Clarendon Press, 2001a.
- Sider, Theodore. "Hirsch's Attack on Ontologese." *Nous* 48, no. 3 (2014).
- Sider, Theodore. "Ontological Realism." In *Metametaphysics*, edited by David Chalmers, David Manley and Ryan Wasserman. Oxford: Clarendon Press, 2009.
- Skyrms, Brian. "Tractarian Nominalism." *Philosophical Studies* 40 (1981).
- Soames, Scott. "Ontology, Analyticity, and Meaning." In *Metametaphysics*, edited by David Chalmers, David Manley and Ryan Wasserman. Oxford: Clarendon Press, 2009.

