
Metaontology and Inconsistent Concepts

An Inconsistency Approach to the Concept of Existence

Master's Thesis submitted for the
Research Master's in Philosophy

Utrecht University

June 10, 2018

Sofia Ellinor Bokros

Student number: 6032966

Supervisor: Prof. Dr. Daniel Cohnitz

Second reader: Dr. Carlo Nicolai

Third reader: Dr. Chris Meyns



Universiteit Utrecht

ABSTRACT

Recent years has seen a surge of new interest and research into neo-Carnapian and deflationary approaches to ontology; it has become increasingly viable to maintain that ontological debates are merely verbal (Hirsch 2011), thoroughly trivial or metalinguistic (Thomasson 2015; 2017), or have no objective or determinate answers (Chalmers 2009). The deflationary stance is in stark opposition to the self-conception of ambitious ontologists, who envision ontological inquiry as the attempt to uncover the fundamental structure of reality (Sider 2009; 2011). In this thesis I argue that the widespread disagreement about the nature of ontology can be better made sense of on the hypothesis that our ordinary concept of existence is defective. Drawing on Scharp's (2013) theory of inconsistent concepts, I propose that our pre-theoretical concept of existence is, in fact, inconsistent: its application conditions are such that in some cases, we can correctly judge that the concept both applies and disapplies to a given object. This has the implication that our ordinary concept of existence fails to reliably track a unique property or structure in the world, as there is no corresponding property or structure which could be both instantiated and uninstantiated at the same time. Consequently, ontological debates are frequently defective as they fail to be about a determinate or substantive subject matter. Nonetheless, instead of deflating ontological debates, I argue that the insight that the concept of existence is defective does not undermine the possibility of substantive ontological inquiry; we can use conceptual engineering to revise or replace our pre-theoretical concept of existence, and then re-formulate the ontological questions of interest. (262 words)

Total word count (excluding abstract): 34 416

TABLE OF CONTENTS

Introduction	1
0.1 Is There Something Wrong with Ontological Debates?	1
0.2 Neo-Carnapian Dismissivism	3
0.3 Dismissing the Dismissivists?	5
0.4 The Inconsistency Approach	7
0.5 Looking Ahead	10
1 Rival Metaontologies	12
1.1 Hirsch's Verbal Disputes View	12
1.2 Thomasson's Easy Ontology	24
1.3 Chalmers' Ontological Anti-Realism	33
1.4 Conclusion	41
2 Inconsistent Concepts	43
2.1 What Are Concepts?	43
2.2 What Are Inconsistent Concepts?	45
2.3 Possessing Inconsistent Concepts	49
2.4 Error Theory and Externalism	56
3 The Inconsistency Approach to Existence	61
3.1 The Problem of Non-Being	61
3.2 Rules of Application for 'Exists'	63
3.3 Metaontological Advantages	66
3.4 Objections to the Inconsistency Approach	70
3.5 Conclusion	77
Concluding Remarks	78
Bibliography	81

INTRODUCTION

0.1 *Is There Something Wrong with Ontological Debates?*

It's not uncommon for philosophers to have the intuition that something is amiss with ontological debates. It's not difficult to see where they are coming from; is it really worth arguing about whether ordinary objects such as tables and chairs *really* exist, or whether the Eiffel tower and Plato's nose together compose an object? Is there even anything here to disagree about in the first place?

In particular to those outside of metaphysics, it seems tempting to dismiss these issues as futile, or of no practical relevance, even if there is nothing wrong with these questions in principle. For instance, it seems to make very little difference to our daily lives if there is a lump of clay there in addition to the statue, or a table in addition to particles arranged table-wise -- regardless of whether these are perfectly good questions in themselves. However, leaving aside the issue of practical relevance which concerns philosophers at large, ontologists themselves are often troubled by uncomfortable questions about the nature of their inquiry. Even subsequent to deciding that ontological puzzles are worthy of attention, a worry remains among metaphysicians that ontological disagreements do not run terribly deep. Is there really a fact of the matter to be settled in all ontological debates, or are ontologists sometimes just arguing about which words should be used to describe the same facts? Perhaps, as it has often been suggested, there aren't always substantive disagreements between the parties to ontological debates; there are simply many, equally good ways to talk.

As an illustration, let's imagine a world which contains exactly two material simples. How many things are there in this world? The nihilist would say two things, and no more, and the universalist would say three, and no less. Now, what are the nihilist and the universalist disagreeing about?¹

At first blush, the answer is of course that they are disagreeing about *the number* of things: Is it two, or is it three? However, granted that the nihilist and the universalist are both counting

¹ The example is inspired by Dorr (2005).

in ordinary English, and do not disagree about the meaning of either of the words ‘two’ or ‘three’, this answer seems puzzling, since they both purport to be describing the same world.

On a closer inspection, then, they must be disagreeing about the number of *things*, and what qualifies as a “thing”. According to the nihilist, only simples qualify as “things,” and seeing as there are two simples, there are two “things.” According to the universalist, the two simples additionally compose a further “thing”, and consequently, there are three “things”. Hence, it seems as if the nihilist and the universalist are in fact disagreeing on the meaning or application conditions of the word ‘thing’, and each is correct in their own language. The nihilist asserts that “there are two things” in her language, and the universalist asserts that “there are three things” in her language, and these two statements do not contradict each other since ‘thing’ has a different meaning in the mouth of each disputant; they are making two different claims. In other words, that there are two *things_n* is entirely compatible with there being three *things_u*. Alternatively, if it turns out that they are not disagreeing about the meaning of the word ‘thing’ after all, it could instead be the case that they are disagreeing about what it means for the thing to really “be there” in the world, or what conditions must be met for two things to “compose” a further thing (in both cases, the nihilist then puts the bar very high and the universalist puts the bar very low). In any event, it appears as if the dispute could potentially be entirely reduced to language, and that there isn’t really much disagreement here after all. The example above would be a prime example of what philosophers call a *verbal dispute*.

The puzzle is of course that, if the above answer is roughly correct, and the participants to the debate are merely disagreeing about language, why do ontological debates like the one above persist? How could it be possible that philosophers have argued over some ontological matters for over two thousand years, and in the end all it comes down to is the meaning of a few words -- should this not have been realised sooner? Too much cynicism thus runs the risk of too easily brushing aside the history of philosophy.

Questions about the nature of ontological debates, such as the one in the toy example above, are central to metaontology. Metaontology asks what it is we are up to when we engage with ontological issues and how we come to know ontological facts. Is ontology a hopeless pursuit? What are the methods that ontology relies on? And do these methods aid us in uncovering deep truths about the structure of reality, or merely lead us into confusions about the linguistic details of our conceptual scheme?

0.2 Neo-Carnapian Dismissivism

The idea that something is seriously amiss with ontological debates has become increasingly prevalent in the recent metaontological literature. Following Bennett (2009), I will refer to this view -- that ontological debates ought to be dismissed -- as *dismissivism*. Dismissivism comes in a variety of flavours. Each variety comes with its own diagnosis of what, exactly, has gone wrong in paradigmatic ontological debates. By far the most popular type of dismissivism is ontological *deflationism* which is the attempt to deflate ontological debates by showing that the participants aren't really disagreeing about the deep matters that they thought they were, or perhaps not disagreeing at all.² Advocates of deflationary views commonly embrace some form of *semanticism*. Semanticist views hold that the source of the trouble with ontological disputes is first and foremost semantic; i.e. about the meanings of the words that are used in ontology, or perhaps even the concept of existence itself.

A commonly embraced deflationary and semanticist view is to say that although there is nothing defective about our ordinary concept of existence, philosophers tend to use it in a problematic way. According to these deflationists, ontologists specifically tend to use the concept of existence illegitimately, which results in ontological debates being defective or senseless. Ordinary existence discourse, on the other hand, is non-defective and wholly functional; when in an ordinary conversation we ask questions about whether there are ghosts, or whether King Arthur ever existed, our existence concept works just fine. The fact that the concept of existence is successfully used in ordinary discourse contributes to the illusion that the philosophers' questions make sense, when in fact they don't.

A view of this kind can be dated back to Carnap (1950), who makes a distinction between internal and external questions. According to Carnap, *internal* questions are questions asked within a linguistic framework, in which the terms used to phrase the questions have determinate meanings and application conditions. Internal existence questions are therefore unproblematic and always straightforwardly answerable either by conceptual analysis or empirical methods; properties, numbers and ordinary objects all exist, so long as these terms can be correctly and meaningfully applied. Ontologists, on the other hand, are trying to ask an *external* existence question -- a question that is external to the linguistic framework -- on which the application conditions are not clearly specified. Accordingly, the only meaningful way of asking

² See Bennett (2009) for a different type of *epistemicist* dismissivism.

external questions is when they are posed as pragmatic questions about which linguistic framework to adopt; i.e. whether talk of “existence”, “numbers” or “objects” is a useful way to talk. External-ontological questions are on Carnap’s view therefore confused pseudo-questions, since they purport to ask something about the world when in fact they can only be made sense of as practical questions about language.

Many contemporary philosophers defend neo-Carnapian views of this kind. Thomasson (2007; 2009; 2015; 2017) argues for a view in just this spirit, which she refers to as *easy ontology*. According to easy ontology, ontological questions can all be easily answered either by conceptual analysis or empirical methods -- as long as we don’t get confused, and stick to using the existence-concept in its ordinary sense. Hirsch (2002a; 2002b; 2009; 2011) is also sympathetic to the idea that the problem lies with how philosophers use quantificational language. He defends a view which he calls *quantifier variantism*, which is the view that the meaning of the quantifier and/or the existence-concept varies in the mouths of disputants of ontological debates. This results in many ontological disputes coming out as merely verbal disputes, just like the toy example given previously, where the disputants are merely talking past each other because they mean different things when they assert that e.g. numbers or properties exist or do not exist.

Another important distinction which can be made among metaontological views is between realist and anti-realist views. For example, Thomasson’s and Hirsch’s views both come out as realist, since according to both of their views, there is a fact of the matter of whether numbers, properties, and so on, exist. On both of their views, the concept of existence has a determinate meaning, and in most cases one of the parties to the ontological dispute has got the meaning right (at the very least, there is a right answer out there) and the other disputants are simply wrong about the use of quantificational expressions. Consequently, there still are ontological truths to uncover -- although uncovering them isn’t necessarily particularly difficult, nor are the truths terribly deep.

On the other side of the realism divide, Chalmers (2009) has argued for anti-realism about ontological debates, which is the thesis that there are no objectively or determinately true or false existence-facts. Chalmers suggests that the meaning of what he terms *the absolute existential quantifier* is indeterminate, which has the result that the meanings of ontological claims also come out as indeterminate. In turn, the indeterminacy of meaning leads to ontological claims frequently lacking determinate truth-values, which makes them objectively or determinately neither true nor false. According to the deflationary anti-realist, then, what’s wrong with

ontological debates is that there are no factual answers to the questions which ontologists are asking, because of a defect in the ontologists' concept of existence.

Hence, if the neo-Carnapian dismissivists are to be believed, ontological debates largely come down to confusion about how quantificational language or the concept of existence work. Once we become clear on the specifics of our linguistic framework, we'll either realise that we were asking pseudo-questions, merely talking past each other, or that there are no objective or determinate facts of the matter. Either way, the problems evaporate. The neo-Carnapian bottom line is that we are best off just sticking to our commonsense ontology; philosophers should redirect their attention elsewhere.

0.3 Dismissing the Dismissivists?

There are however some reasons to suspect that the neo-Carnapian rationale cannot be the whole story. For example, many ontological disputes and questions seem to persist or reincarnate even though it has been pointed out by many that they are "merely verbal". If all ontological puzzles only came down to confusions about language, then surely all the disputants should sooner or later have come to terms with this fact and directed their attention elsewhere?

Perhaps philosophers are merely slow at catching on, but the most natural explanation for the continued interest in ontological questions would appear to be that there are, in fact, interesting and non-trivial ontological problems to be addressed.

In spite of the barrage of criticism, there are thus still those who continue to defend the conviction that existence-questions are deep and difficult metaphysical questions which are not to be dismissed lightly. For example, Sider (2009; 2011) and van Inwagen (1998; 2009) are among those who have defended the view that ontological questions are meaningful and rigorous philosophical questions.

In response to the dismissivists, Sider (2009; 2011) has argued that the deflationary arguments needn't trouble the serious ontologists. Even if the deflationists turn out to be correct about our existence-concept as we use it in ordinary conversational English, the ontologists have a plan B: *Ontologese*. Ontologese is a special technical language introduced specifically for doing ontology, in which we can stipulate that the existential quantifier is to refer to reality's ontological structure. By means of this language, we can simply restate the questions we were trying to ask; the existential quantifier -- at least how it is used in the ontology room -- refers to reality's

ontological structure, and attempts to carve nature at its joints. The task of ontology is accordingly to uncover the structure of reality. This squares better with how the ontological project has traditionally been conceived, and aims to reinvigorate the age-old questions.

I think there are insights to be gained from both the dismissivist and the anti-dismissivist positions. To see how this could be possible, we first need to disambiguate the central metaontological question which both sides are trying to answer to. The central question of metaontology is “What are ontological disputes about?” We can distinguish at least three different senses of this question:

1. *Modal*: What is it *possible* for ontological disputes to be about?
2. *Descriptive*: What are current ontological disputes *actually* about?
3. *Normative*: What *should* ontological disputes be about?

Once we distinguish these three senses, we can see that different answers can be given to each question. In contemporary metaontological debates, the three senses are rarely explicitly distinguished, and sometimes even conflated. To some extent, this conflation is innocuous, as the answers given to the possibility question do constrain what answers can plausibly be given to the descriptive and normative questions. But disambiguating the three senses of the metaontological question can help us make room for a more nuanced middle ground, drawing on the insights provided by both the dismissivists and the anti-dismissivists.

What should now become apparent is that the deflationists typically focus their analysis on the descriptive question, trying to analyse what current and traditional ontological debates are about. The non-deflationary realists, on the other hand, tend to privilege the normative question; although they may admit that some current ontological debates are potentially defective, this is for them no reason to stop doing ontology, but rather to do ontology better.

Plausibly this comes down to a tacit disagreement, or divergent assumptions, about the possibility of doing ontology in the first place. Dismissivists are inclined to think that doing substantial ontology isn’t really possible -- either due to there not being any objective or determinate answers, or because there is no way of asking meaningful and coherent ontological questions that aren’t trivial. Ontological questions are either so easily answered that there is no need to spend time on them, or they are unanswerable pseudo-problems.

Anti-dismissivists are more optimistic about the possibility of doing ontology, and typically think that there are meaningful philosophical questions about “what exists” to be asked, regardless of whether our ordinary language is up to the task. As Sider succinctly articulates the

sentiment; there's no detour to be taken around the entirety of fundamental metaphysics (2009, p. 420).

In this thesis I will attempt to assess the deflationary accounts, and the import of their arguments for the nature and possibility of doing ontology. I will defend the view that contrary to the deflationists' conviction, ontology cannot be dismissed as easily as they make it seem. In part, this is because no cogent arguments appear to have been given in support of the impossibility of doing ontology. Deflationists typically spill most of their ink on trying to show that our actual ontological debates are defective or trivial (i.e. on the descriptive issue) -- but this does very little to show that doing ontology well is *impossible*. If our commonsense concepts aren't up to the task of asking coherent or non-trivial questions; who's to say that new concepts cannot successfully be introduced?

Nevertheless, I will leave the modal and normative issues to one side for now, and first focus on the descriptive question. The main focus of this thesis will be on the point which it seems that the deflationists do get roughly right; current and traditional ontological debates do indeed seem to frequently be defective. In this thesis I will attempt to defend a different approach to diagnosing what has gone wrong in ontological debates; namely, the view that existence is an inconsistent concept.

0.4 *The Inconsistency Approach*³

This thesis examines whether we can make more sense of what is going on in ontological debates on the hypothesis that existence is a defective concept. More specifically, I will attempt to defend the idea that existence is an *inconsistent* concept. The very rough idea behind the theory of inconsistent concepts is that some concepts have incoherent rules of application, which create semantic deficiencies and philosophical paradox.

To elaborate, every concept has its own set of rules of application, or application conditions, which dictate how that concept may be correctly used in natural language. All competent speakers of a language have implicit knowledge of these rules, and intuitively follow them in using the concept and its corresponding terms. However, because natural language and its associated conceptual framework is an organically evolved human construct, these rules of

³ Following standard conventions, in this thesis I use SMALL CAPS to denote concepts (i.e. BLUE is shorthand for referring to the concept of blueness), single quotation marks (") for mention, double quotation marks (") for use and simultaneous use and mention, and *italics* for emphasis or to introduce technical terms.

application form far from a perfectly complete and consistent system. Accordingly, concepts can be intensionally defective in many ways; they can for example be unsatisfiable, partial, vague, empty or inconsistent. When a concept is inconsistent, its rules of application are such that they lead us into accepting inconsistencies; they make us inclined to make contradictory judgments and to apply the concept in inconsistent ways. In other words, they produce paradox.

Perhaps the most well-known inconsistent concept is the concept of truth (Scharp 2007; 2013; Patterson 2009; Azzouni 2007; Eklund 2002).⁴ The concept of truth is constituted by two principal rules of application, famously known as the T-schema (where ‘S’ stands for a sentence):

(T-In) If S, then ‘S’ is true.

(T-Out) If ‘S’ is true, then S.

Although both of these rules seem intuitively correct and individually unproblematic, they create an inconsistency when applied to certain problematic sentences, such as sentence ‘L’:

L ‘L’ is not true.

Sentence ‘L’ allows us to formulate the *Liar Paradox*. First, assume that ‘L’ is true. If so, then ‘‘L’ is not true’ is true, and by rule (T-Out), we can infer that ‘L’ is not true. But that’s a contradiction. So by reductio, ‘L’ can’t be true. Therefore, it must be the case that ‘L’ is not true. But then by (T-In), we can infer that ‘‘L’ is not true’ is true, so (substituting ‘L’ for ‘‘L’ is not true’) ‘L’ is true. Again, we have a contradiction. Hence, if we accept the T-schema, assuming either that ‘L’ is true or that ‘L’ is not true leads us into contradiction. This means that anyone who possesses and is competent with our concept of truth (as embodied by the rules of application spelled out in the T-schema), is going to be inclined to accept a contradiction when confronted by a sentence such as ‘L’. In other words, if the rules of application for the concept of truth are given by the T-schema, then TRUTH is an inconsistent concept.

The inconsistency approach to existence holds that something very similar is going on in the case of EXISTENCE as is going on in the case of TRUTH. That is, our ordinary concept of existence has rules of application that tempt us into accepting contradictions, which create paradox (the most well-known paradox being that posed by negative existentials or *the problem of non-being*). In addition, the inconsistent rules of application give rise to semantic confusions and

⁴ There are of course many philosophers who do not regard truth as an inconsistent concept. This is not intended as an argument for the view that truth is an inconsistent concept, but merely as an illustration of the underlying idea.

indeterminacies, which results in ontological debates sometimes being defective in just the ways that the deflationists maintain: due to conceptual defects, ontological debates sometimes end up being merely verbal, metalinguistic or insubstantive. In addition, because the concept of existence has inconsistent rules of application, it can sometimes be the case that when one ontologist says that “numbers exist” and the other retorts that “numbers don’t exist”, both of their utterances are in fact correct applications of the very same ordinary concept of existence.

I think this view has many virtues, and better explains why ontological disputes have persisted for so long -- but also why they arose in the first place. Because the approach locates the culprit in the concept of existence itself, the approach to metaontology is semanticist, although neither deflationary nor anti-realist. It is also not dismissivist, as I don’t believe the observation that ontological debates are defective warrants the attitude that we ought to dismiss them; the answer should in the first instance rather be to try and correct them. If a defect in our concept of existence is responsible for why ontological debates are defective, this need not show that ontological questions are confused pseudo-questions, or that substantive ontological inquiry is impossible -- but merely that the way in which we are currently doing ontology generates unnecessary metalinguistic confusions. The inconsistency approach suggests a natural way out of this impasse; via conceptual engineering.

Some philosophers have hypothesised that traditional philosophy is, for a large part, the study of inconsistent concepts (see Scharp forthcoming). Concepts such as truth, knowledge, freedom, personal identity, rationality, justice, goodness, time and causation are potentially all inconsistent or otherwise defective concepts, which explains why we get ourselves into paradox when we try to analyse these concepts. However, we needn’t necessarily subscribe to such a radical metaphilosophical view to recognise that conceptual defects are worth taking seriously as culprits or contributing factors to specific philosophical problems, confusions and disagreements. Finding and understanding these conceptual defects might therefore be a fruitful way to make progress in some philosophical debates.

The aims of this thesis are thus twofold. The primary aim of the thesis is to examine whether understanding the concept of existence as an inconsistent or otherwise defective concept can help us make sense of ontological debates, as outlined above.

The secondary aim of the thesis is to figure as a metaphilosophical test case. I would like to view the project of this thesis partly as a case study for testing the broader metaphilosophical hypothesis that philosophy is generally concerned with inconsistent or otherwise defective

concepts. If the inconsistency approach to existence proves to be a coherent and fruitful approach to making sense of ontological debates, then possibly the inconsistency approach is also a fruitful approach to take to other philosophical concepts and topics. Thus, regardless of whether the inconsistency approach turns out to be the correct approach to ontology, perhaps it can at the very least serve as a methodological experiment; in the best case as a model.

0.5 Looking Ahead

I will conclude this introduction with a brief summary of what's to come. In the next chapter, I will return to some of the rival views mentioned above and discuss them in more detail, to show exactly where I think they go wrong. I will discuss the deflationary views of Hirsch, Thomasson and Chalmers briefly outlined above, and argue that none of the views succeed in deflating ontological debates. In addition, neither the interpretation of ontological debates as merely verbal disputes or as metalinguistic negotiations provide a sufficiently charitable or explanatorily adequate interpretation of what it is that ontologists are doing when they engage in an ontological dispute. That is, they both fail to explain all the phenomena in a plausible or charitable manner.

In chapter two, I make a detour into the theory of inconsistent concepts. The discussion is primarily based on Scharp's (2013) account of inconsistent concepts and concept possession. I will criticise his account of inconsistent concepts and propose some amendments. The aim is to make plausible that the idea of inconsistent concepts is a coherent one -- and one with great explanatory potential at that. In addition, I will answer some very recent objections to the idea that there are inconsistent concepts due to Herman Cappelen (2018).

Chapter three makes use of the theory of inconsistent concepts outlined in chapter two and applies it to the concept of existence. I will argue that the problem of non-being provides us with at least *prima facie* grounds for analysing EXISTENCE as an inconsistent concept, and that the view has great explanatory potential as the basis of a metaontological account. The aim is not to defend the specific analysis *per se*, but the idea that the predicate contains an inconsistency; the overall argument takes the form of an inference to the best explanation. I will argue that the inconsistency approach has many explanatory advantages over the metaontological accounts discussed in chapter one, while still preserving many of their insights. Finally, I answer some

objections to the inconsistency approach and briefly compare it to some suggestions made by Theodore Sider's ontological realism (2009; 2011).

Lastly, I conclude the thesis with some final remarks on the project of ontology as I view it, and some ideas on how we can use the findings to make progress in ontological debates.

1.

RIVAL METAONTOLOGIES

In recent years, deflationary views of ontology have gained increased popularity in the metaontological literature (see e.g. Chalmers, Manley & Wasserman 2009; Blatti & Lapointe 2016). What these metaontological accounts have in common is that they all attempt to deflate ontological debates; to show that they are not as deep, significant or substantive as the serious ontologists believe them to be. A common sentiment is that ontological debates can be dismissed as fruitless; either because they are entirely misguided, there is little at stake, or because there are no objective or determinate answers. This chapter will examine three prominent such deflationary views; the *Verbal Disputes View* due to Eli Hirsch, Amie Thomasson's *Easy Ontology* and David Chalmers' *Ontological Anti-Realism*. Although their metaontological outlooks are similar, these are actually very different object-level views of ontological disputes. In fact, each of these three views have their own analysis of what is happening at the semantic level of ontological disputes, which is then paired with a corresponding deflationary metaontological diagnosis of the nature and possibility of doing ontology. I will begin by discussing Hirsch's verbal disputes view along with quantifier variantism, and secondly I will discuss Thomasson's easy ontology and her interpretation of ontological debates as metalinguistic negotiations. Lastly, I will turn to Chalmers' anti-realism. I will argue that neither of these three accounts provide an entirely satisfactory explanation of what goes on in ontological debates. And perhaps most significantly, even if either of the three accounts were assumed to be correct, none of the accounts decisively deflate ontological debates, as none of them provide conclusive reasons for thinking that substantive ontology is an impossible project.

1.1 Hirsch's *Verbal Disputes View*

1.1.1 *Mere Verbal Disputes and Quantifier Variantism*

Perhaps the most widely discussed deflationary metaontological view is the verbal disputes view, primarily associated with the work of Eli Hirsch (2002a; 2002b; 2005; 2009; 2011). On this view, ontological disputes are not genuine disputes over worldly matters of fact, but arise as a result of a tacit disagreement over meaning. Accordingly, once the metalinguistic disagreement is resolved, the ontological dispute evaporates.

An illustrative example of a verbal dispute given in the literature is that of two people arguing over whether an alcoholic cocktail made with apple liqueur, served in a tall V-shaped glass, is a martini (Bennett 2009). One of the disputants claims that she is drinking a martini, while the other disputant emphatically denies her claim; “*That* is not a martini!” Clearly what is going on in such a dispute is not a disagreement over the facts; i.e. it is not the case that the disputants have different perceptions or beliefs about the contents of the drink that is on the table. Rather, the disagreement can be accounted for wholly in terms of a disagreement over the meaning of the word ‘martini’. According to the purist, only a mixture of gin, vermouth, and perhaps an olive truly qualifies as a martini; according to the casual customer, pretty much any alcoholic cocktail served in the classic V-shaped glass can be said to be a type of martini. The dispute is therefore “merely verbal” in the sense that it is entirely reducible to a metalinguistic disagreement about how the word ‘martini’ is to be defined; both parties could explicitly agree on the non-linguistic fact that the drink in question consists of vodka, apple juice and apple liqueur, and is served in a V-shaped glass.

More precisely, a verbal dispute can be characterised as a dispute in which each participant can be interpreted as speaking the truth in their own language (Hirsch 2009; 2011). For example, in the disagreement over the martini, if one of the disputants is understood as asserting that she is drinking an alcoholic cocktail in V-shaped glass, and the other disputant is understood as asserting that the drink in question is not made up of gin and vermouth, they are clearly both asserting something that is true. Due to the difference in how they use the word ‘martini’, the disputants are most charitably interpreted as making very different claims that are not in fact in conflict with each other. If each of the disputants had full insight into how the other party uses the word, they could in principle acknowledge that the other party is asserting a true claim. Hirsch thus defines a verbal dispute as follows:

[A verbal dispute] is a dispute in which, given the correct view of linguistic interpretation, each party will agree that the other party speaks the truth in its own language. This can be

put more briefly by saying that in a verbal dispute each party ought to agree that the other party speaks the truth in its own language.

(Hirsch 2009, p. 239)

On Hirsch's view, some metaphysical disputes are merely verbal disputes. Specifically, some metaphysical disputes are such that each party in fact agrees on all the non-linguistic facts, but the dispute arises due to differences in how each party uses language.⁵ If these differences in meaning were to be explicitly spelled out, each party could agree that the other speaks the truth in their own language, and the dispute would evaporate.

In the ontological case, this amounts to the thesis of *quantifier variantism*. Quantifier variantism holds that the meaning of quantificational expressions varies in the mouths of the disputants. As a result, existence claims such as "tables exist", can have several different meanings depending on who asserts the claim and in what context. According to quantifier variantism, there is thus no privileged sense of the quantifier at work in existential claims, but instead there are "many possible perspectives on 'the existence of objects', which all are adequate for describing the same facts, the same 'way the world is'" (Hirsch 2004, p. 231).

Quantifier variantism is primarily motivated by the principle of charity. The principle of charity states that in interpreting a speaker one ought to interpret the meanings of their expressions such that their claims insofar as possible come out as true or, minimally, not completely unreasonable.

Central to linguistic interpretation is the presumption that the correct interpretation is the one that makes people's use of language as reasonable as possible. In interpreting a language, there is therefore an overwhelming, if in principle defeasible, presumption that typical speakers make perceptual assertions that are reasonably accurate, and that they do not assert relatively simple sentences that are a priori false.

(Hirsch 2009, p. 240)

In order to respect these considerations, then, we should try to find the interpretation which makes the other party come across as saying something that isn't obviously incorrect. For instance, when the casual customer from the example above asserts "This is a great martini", the principle of charity recommends that we shouldn't ascribe to them the false belief that they are sampling a drink made of gin and vermouth, but rather interpret them as if they mean something

⁵ To clarify, Hirsch does not make the blanket claim that all metaphysical or ontological disputes are verbal, merely some debates. For example, Hirsch (2009) argues that the debate between *Endurantists* and *Perdurantists* is merely verbal, and he has claimed elsewhere (2011) that e.g. the debate between nominalists and platonists is an example of a debate which is not verbal.

else by ‘martini’. Similarly, in the case of ontological disputes, where for instance the nihilist claims that “tables don’t exist”, while the universalist claims that they do (and neither appears to be making some kind of empirical mistake), it therefore seems that in order to be as charitable as possible to each side of the dispute, we ought to interpret the nihilist and the universalist as meaning different things when they make ontological claims. When the nihilist claims that “tables don’t exist”, we should not interpret the nihilist as claiming that they have never seen or used a table, or something else that is obviously false. Instead, we should conclude that they must have something else in mind with their assertion. This means that in the respective languages of the nihilist and the universalist, existence claims and other quantificational expressions must vary in meaning such that each disputant can charitably be interpreted as saying something that is true in their own language.

The externalist about meaning might object to the idea of quantifier variance, and the occurrence of verbal disputes more generally, by maintaining that the meaning of expressions is determined by the linguistic community. According to this externalist argument, regardless of what the disputants *think* they mean, their terms have the same meaning in virtue of the disputants belonging to the same linguistic community. However, this is besides the point, as we can just make a distinction between the literal meaning of the term and the speaker’s meaning, and the dispute remains “merely verbal” in virtue of at least one of the disputants being mistaken about the literal meaning. Hirsch suggests that we simply stipulate, for the sake of discussion, that each disputant forms its own linguistic community, such that the meaning of the terms in her language is determined precisely by her own use, or the speaker’s meaning. We can then still characterise a verbal dispute as one in which each disputant ought to agree that the other is speaking the truth in their respective idiolect.

Likewise, the verbal disputes view is not threatened by the fact that in some verbal disputes, one of the disputants is obviously right and the other is wrong. The fact that one of the disputants is using the term correctly and that the other is misusing the term does not make their dispute any less verbal, as the disagreement is still reducible to a disagreement over the meaning of words. For example, in the disagreement over the “martini”, the purist might turn out to be strictly speaking correct, all the while the casual customer can still be thought of as speaking the truth in her own language. Hence, even though it is possible to win a merely verbal dispute, it remains a rather trivial victory.

Consequently, if we accept that some ontological disputes are merely verbal, there appears to be very little at stake in these disputes. At best, it seems that these disputes amount to a matter of linguistic correctness or pragmatic linguistic choice. On Hirsch's view, provided that everyone intends to be making their philosophical claims in our common language, the only issue to be settled is then which of the disputants is speaking ordinary English; that is the only "ontological truth" which we can hope to uncover (Hirsch 2011, p. 182). Just like the disagreement over the martini would be settled by establishing that in ordinary English the drink does not qualify as a martini, we can settle ontological disputes by establishing which ontological claims are correct in ordinary English. This has the result that the correct ontological view must be aligned with our commonsense ontological judgments; if most ordinary speakers of English would agree that tables exist, then the nihilist or eliminativist ontological view is plainly false. Hence, Hirsch does not actually himself accept that the quantifier varies in meaning in ordinary English -- although he does admit that there are several different equally good quantifier meanings that could be adopted by those who are willing to deviate from standard use.⁶

In summary, if the verbal disputes view of ontological debates is correct, it would appear that ontological debates are highly shallow. There is no non-linguistic fact at issue in these debates; instead ontological questions reduce to questions about how ordinary English works, and how our language is correctly to be used -- questions which are perhaps best left to the linguists.

1.1.2 Objections to Quantifier Variantism

Many critics have raised serious doubts about the possibility of quantifier variance, and whether we can make sense of the idea that there are multiple existence-like concepts with slightly different meanings. If quantificational expressions cannot vary in meaning, that makes it more difficult to maintain that ontological disputes are merely verbal (although not impossible, as instead one could maintain that words such as 'number', 'table', or 'object' vary in meaning). Moreover, quantifier variance is bound up with linguistic pluralism, i.e. the thesis that there are several possible equally good languages for doing ontology. In this section I will deal with the main objections against quantifier variantism and linguistic pluralism from anti-realism and

⁶ Hirsch thus allows for the possibility that "revisionist" metaphysicians and ontologists may introduce their own technical language, but in that case "they have to say that, and not pretend that they are expressing in plain English a substantive and controversial philosophical discovery" (2002b, p. 106).

linguistic relativism, as well as Matti Eklund's *semantic argument* (2009). In addition to these objections, Sider (2009; 2011) has also raised some problems for quantifier variantism on the basis of externalist considerations of reference magnetism and Lewisian naturalness. I will nevertheless leave Sider's arguments to one side for now, and return to Sider's views in chapter three when I discuss objections to the inconsistency approach.

The first and perhaps most common objection to quantifier variantism is the objection from anti-realism. The anti-realism objection states that quantifier variantism reduces ontological issues to linguistic choice, and thus results in anti-realism or relativism. In other words, the objection says that, if the quantifier varies in meaning, what exists either depends entirely on which language we happen to speak, or is relative to our linguistic framework. This objection, however, is simply a failure to understand the claim that quantifier variantism makes; quantifier variantism, properly understood, does not imply any kind of anti-realism or linguistic relativism. What quantifier variantism maintains, is that the meaning of 'exists' and other quantificational expressions is a matter of linguistic choice. But once we have made a choice and assigned a semantics to the quantifier, what "exists" is strictly dependent on the world. To use an analogy, which colour is denoted by the expression 'green' is trivially dependent on linguistic choice; we could use the word 'green' to mean blue or to mean red if we wanted to -- but as it so happens, the word 'green' means green. Accordingly, it is indeed a matter of linguistic choice whether grass is 'green' or the sky is 'green', but that does not imply that grass being green is a fact determined by or relative to language. The objection that quantifier variance implies anti-realism or linguistic relativism therefore rests on a use-mention fallacy (see Hirsch 2002a for a sustained defense of the point that quantifier variance and realism are compatible).

A related, and more serious worry is that, although it seems trivially true that we could use 'exists' to express a number of different meanings or concepts, these different meanings could not all be quantifier expressions. It is of course trivially true that we could use 'exists' to express the concept of swimming (i.e. the meaning of the word 'swims'), and so by saying that 'unicorns don't exist' we would be saying that "unicorns don't swim". However, swimming is not an existence-like concept, although we could certainly express it by using the word 'exists'. Hence, for an expression to be a quantifier expression, it needs to have a particular *kind* of meaning, and this meaning cannot differ significantly from the concept of existence in ordinary English. Therefore, the objection goes, the quantifier could only have one meaning, and if that meaning was different, it wouldn't be a quantificational expression. As such, the worry is that

there simply couldn't be multiple quantifier-like or existence-like expressions, as we would be changing the subject.

In response, it should be granted that it is undoubtedly true that not just any concept qualifies as existence-like. But given that there are several conceptual analyses of existence present in the philosophical literature, it seems rather plausible that there should be several good candidate concepts, embodied by these different analyses, for the concept of existence. Even if turns out that one of these conceptual analyses are correct and that the others are mistaken, the mistaken conceptual analyses seem to still express something that is *existence-like*. So it seems difficult to insist that it is impossible for there to be more than one existence-like concept (even if there is maybe only one concept of existence in ordinary English). Exactly where we draw the line between a concept that is sufficiently existence-like to qualify as a quantificational expression is nonetheless up for debate -- but we need not settle that issue here. A further point in response to the objection is that one ought to be careful not to confuse the quantifier as a logical operator, '∃', with quantificational expressions in ordinary English, or the concept of existence in natural language. In logic, it is certainly true that '∃' has a formal, fixed rule of use. That being said, I see no *prima facie* reasons for why there couldn't be several different candidate meanings for quantificational expressions or the concept of existence in ordinary English.

As already noted above, quantifier variance implies *linguistic pluralism*, i.e. the thesis that there are several equally good languages in which we can do ontology. (Note, however, that linguistic pluralism could be true without quantifier variance being true.) According to some writers, such as Eklund (2009), linguistic pluralism also implies a kind of ontological pluralism. The thought is that if there are several different existence-like expressions, or several languages in which one can describe the world equally correctly, this appears to imply that in some sense there are several correct ontologies, each applying their own notion of existence. However, it is of crucial importance not to conflate this type of ontological pluralism with linguistic relativism or anti-realism mentioned above. Linguistic pluralism allows that, if our concept of existence varies in meaning, it can be true in one language that 'Numbers exist', and false in another. So in some sense, both Platonism and Nominalism can come out as correct, depending on which language (or more accurately, which existence-like concept) we employ. Crucially, these different languages must use different notions of existence, such that 'Numbers exist' means different things in the nominalist and platonist language -- the sentence 'Numbers exist' must express two different propositions; "Numbers exist_p" in the platonist language, and "Numbers exist_N" in the

nominalist language. This means that, in fact, it would be incorrect to translate ‘Numbers exist_p’ in the platonist language into ‘Numbers exist_N’ in the nominalist language, because these sentences mean slightly different things. Hence, the platonist and the nominalist do not have to disagree on what the worldly facts are, or what the world is like; they merely describe the world in different ways, using different concepts.

Eklund (2009) has objected to this kind of linguistic pluralism (he calls it ontological pluralism) with an argument which he calls *the semantic argument*. Consider two sentences in two different languages, S_p in the Platonist language L_p , which states that “The number 2 exists_p”, which is true by the pluralist’s lights, and S_N in the Nominalist language L_N , which states that “The number 2 exists_N”, but which is false. Eklund writes:

But what should the *nominalist* say about the truth-value of the sentence S_p of L_p ? It seems that by the pluralist’s lights she should say it’s *true*. Why shouldn’t she? The nominalist, like everyone else, can recognize the truth of ontological pluralism, and recognize that there are some platonistic languages out there, even if hers is not one of them. But mustn’t the singular term refer for the sentence to be true? But then, so the objection to pluralism that I want to consider goes, the nominalist must concede defeat! For then it can be concluded, in L_N , that “2” refers, and that there are numbers.

(Eklund 2009, p. 145, emphasis in original)

However, it appears to me that this objection begs the question against the pluralist. Firstly, note that we could consider two other sentences, S_N^* and S_p^* , in L_N and L_p respectively, which both say that “The number 2 does not exist”. In this scenario, S_N^* in L_N is true, while S_p^* in L_p is false. Now, what should the platonist say about S_N^* in L_N ? He should of course admit that it is true! So it is not the case the platonist automatically wins, as the platonist must also admit that there are languages in which it is true to say that “numbers don’t exist.”

As for the issue of reference -- what it appears that Eklund is doing here, is that he is assuming a specific (platonist) notion of existence as correct from the outset. In fact, he appears to be assuming that reference is both necessary and sufficient for existence⁷ -- but this is a very particular and controversial assumption. More pertinently, it begs the question against the quantifier variantist, who of course maintains that there are several candidate notions of existence -- not just the one presupposed by Eklund that closely ties existence to reference. (If it seems implausible that existence and reference can come apart, we can instead acknowledge that there could be several notions of reference, such that even if it is true in L_p that “‘2’ refers_p”, it

⁷ At the very least, he is assuming that is sufficient. But if reference is only sufficient but not necessary for existence then it is unclear why the nominalist should grant that that ‘2’ refers in the platonist language in the first place.

might still be false in L_N that “‘2’ refers_N” -- but crucially, in order to conclude in L_N that “‘2’ exists_N”, it is not sufficient that “‘2’ refers_p”. Alternatively, the pluralist could maintain that ‘2’ refers to something different in L_N and L_p (although this would no longer be quantifier variantism).)

To drive the point home, if we grant Eklund’s claim that reference is necessary and sufficient for existence, it seems that we can argue as follows. If, as we acknowledged above, S_N^* (“The number 2 does not exist_N”) is true in L_N also by the platonist’s lights, and if reference guarantees existence, it cannot be the case that ‘2’ refers. And since ‘2’ does not refer then the platonist must concede, in L_p , that there are no numbers. If reference is necessary and sufficient for existence, and if both parties can genuinely make true existence-claims in their own languages, the argument must work both ways. Eklund’s “semantic argument” therefore fails as he starts from contradictory premises and begs the question against the pluralist.

In summary, none of the objections to quantifier variantism considered thus far have succeeded in refuting the thesis. Quantifier variance therefore seems as of yet to be a coherent possibility. Whether the meaning of the quantifier as it is used by ontologists *actually* varies in this way, however, is a much more difficult (and to some extent an empirical) question. That being said, I think that, even if we grant quantifier variance, there are other, more serious problems with the verbal disputes view which ultimately make it untenable, as will be discussed in the next section.

1.1.3 Objections to “Merely” Verbal Disputes

As should be clear from the previous section, the problem with the verbal disputes view does not lie with its reliance on quantifier variantism. Rather, the genuinely problematic aspect of the view, in my opinion, lies with its metaontological diagnosis of ontological disputes as “merely verbal.” There are two main problems with this metaontological outlook. Firstly, it renders inexplicable how ontological disputes have managed to persist for such a long time, and secondly, it appears that considerations of charity do not in fact support the verbal disputes view, as although the account is charitable towards what the serious ontologists are *saying*, that comes at a cost of being highly uncharitable to what the serious ontologists are *doing*. I will discuss these and some related objections below. Finally, I will note that, even if ontological disputes turn out to be merely verbal, this does not automatically serve to deflate ontological debates, as the inference from “mere verbalness” to triviality or insubstantiveness does not follow.

The first objection to characterising ontological disputes as merely verbal is that dismissing ontological disputes as verbal fails to adequately explain why it is that ontological disputes have persisted for so long. When was, for example, the last time you came to the realisation that a longstanding disagreement you had with a friend turned out to be merely verbal? If there really was nothing more at stake in these ontological disputes than the correct use of English, it really is quite puzzling why these debates are still ongoing, or why ontological questions continue to interest philosophers. In addition, if ontological disputes were merely verbal, this should realistically have been realised sooner, given that we are normally very adept at detecting disputes which merely turn on the meanings of words. Thus, two disputants arguing about whether Pluto is a planet, or whether anyone who believes in equal rights for men and women is a feminist, would under normal circumstances recognise very quickly that they are assuming different definitions of ‘planet’ or ‘feminist’. As Brendan Balcerak Jackson articulates it:

As conversationalists we are quite used to coping with the possibility that we might be talking past each other, and we are very adept at detecting this and repairing it when it happens in day-to-day conversations. [...] It would be astounding to learn that merely verbal disputes over ‘pain’ or ‘material object’ have somehow managed to survive undetected across decades or even centuries of discussion, despite the continual efforts of the parties to clearly articulate their subject matter and identify their points of disagreement.

(Balcerak Jackson 2014, p. 43).

Naturally, disputes sometimes persist even though the participants come to realise that the dispute is merely verbal. Sometimes people feel very strongly about words, or are unwilling to admit their mistakes. It is not completely far-fetched to imagine an individual who keeps insisting that they are drinking a “martini”, or that they are not a “feminist”, just out of sheer stubbornness or a strong normative conviction that *this is how the word ought to be used*. (This type of normative dispute over how words ought to be used will be discussed in more detail in the next section.) However, in such cases one is usually fully aware of the fact that one is merely disagreeing about the meaning of a word, and not about substantial matters of non-linguistic fact. This observation squares quite badly with the self-conception of serious ontologists who generally do not take themselves to be engaged in verbal disputes at all.

In addition to being unrealistic, then, the verbal disputes view is highly uncharitable to the serious ontologists. In fact, it portrays these ontologists as confused about the very project they are engaging in, as clearly they haven’t yet realised that the only issue to be settled regards the meanings of words in ordinary English. This is completely at odds with the self-perception of

the serious ontologists, as it fails to account for the intuition that the ontologists have that they are inquiring into substantive metaphysical questions. While it might not be *impossible* for the serious ontologists to be mistaken in this way, it would be quite astonishing, considering that philosophers are generally highly reflective in their conduct. Therefore, even though the verbal disputes view succeeds at being charitable towards what the ontologists are *saying*, it is simultaneously highly uncharitable towards what they are *doing* (Horden 2014; Thomasson 2017).

As Horden (2014) has argued, Hirsch's verbal disputes account in fact faces a dilemma when it comes to charitably interpreting what the serious ontologists are up to. Either the ontologists are unwittingly making highly trivial claims in their own respective languages, in which case they are making these claims *unintentionally*, or they are *intentionally* making the same trivial claims. Neither option is particularly attractive.

On the first horn of the dilemma, although these philosophers surely intend to be making non-trivial, philosophically interesting, ontological claims, the verbal disputes view appears to imply that these philosophers are, in fact, unintentionally making highly trivial claims in their own respective languages. But this way of characterising the situation seems self-undermining, since if the ontologists are making claims unintentionally, it isn't clear what ultimately determines the contents of their utterances, as appealing to the disputants' communicative intentions is no longer an option. It then seems that Hirsch would need to appeal to externalist considerations in order to bridge the gap between what the ontologists are intending to communicate and what they, in actuality, unintentionally communicate -- but according to Horden (2014, p. 233), such an appeal would be circular as it would already presuppose that the speaker's respective linguistic communities use language in the way that Hirsch claims.

The remaining option for Hirsch is to take on the second horn of the dilemma, and to claim that the ontologists are making these highly trivial claims in their own respective languages intentionally. But then Hirsch would also require an explanation for why the ontologists, completely knowingly, are making these highly trivial claims. It could of course be the case that the ontologists know full well exactly what they are claiming, without necessarily thinking of it as trivial -- plausibly because there are so many others out there who disagree. (Frequently when things seem trivial or obvious to oneself, the fact that others disagree gives one reason to reconsider whether the claim is really that trivial or obvious.) However, this explanation requires for it to work that the ontologists are unaware of the fact that the other disputants mean

something different with their words, and as noted earlier, this might appear rather implausible, considering that we are normally very adept at detecting linguistic misunderstandings.

In either case, whether the ontologists are making these trivial claims intentionally or unintentionally, they do not, on Hirsch's picture, come out as particularly reasonable. It follows from the dilemma that, on the verbal disputes view, either the ontologists must be mistaken about the contents of their own utterances, or they must be mistaken about the contents of the utterances of the other parties to the dispute. Be that as it may; we can conclude that considerations of charity do not appear to support the verbal disputes view after all.

A final worry worth pointing out is that, considering that both the nihilists and the universalists at least started out as competent speakers of ordinary English, but now speak their own respective languages, a shift in their languages must at some point have taken place. The verbal disputes view fails to convincingly explain why and how this shift has occurred. According to Hirsch, "the philosopher has somehow confused himself into speaking a new language without realizing it" (2002, p. 61). This, however, appears highly counterintuitive, as hopefully we engage in philosophical reflection in order to become less confused; not to *create* confusion. On Hirsch's view it would appear that generating confusion is the only thing which philosophy, or at least ontology, really accomplishes. I am not terribly sympathetic to this outlook, but more relevantly, Hirsch owes us some kind of explanation of how and why philosophical reflection only serves to make us more confused instead of less.

It thus seems rather unlikely that ontological disputes can be written off as merely verbal. Before I conclude the section, however, I briefly wish to note that, even if we grant that ontological disputes are merely verbal, Hirsch's deflationary conclusion does not automatically follow. As an illustration, one might have a merely verbal dispute about whether Pluto is a planet -- yet, this does not make facts about Pluto's planethood insubstantive or trivial. Establishing whether Pluto has sufficient mass to assume hydrostatic equilibrium, or has cleared the neighbourhood around its orbit, are certainly highly non-trivial matters. Hence, the disagreeing ontologists could very well be making substantive, non-trivial claims in their own languages, despite their dispute being merely verbal. To some degree, Hirsch recognises a similar point, as he maintains that his view is realist, and that quantifier variance by no means implies anti-realism, as discussed above. Hirsch thus recognises that the disputants speaking in their own languages are indeed both saying something true; they are simply uttering two different, non-conflicting propositions. It is therefore unclear why Hirsch believes that it should follow from the fact that

ontological disputes are merely verbal that all matters ontological are shallow or without substance.

The missing background assumption appears to be that, if the participants were to speak the same language, there would be nothing left to disagree about, since the facts of ontology are so blindingly obvious or trivial. But it is not self-evident that this is the correct prediction. Interestingly, pointing out that the debate is merely verbal does not seem to have the effect of dissolving disagreement. As Chalmers observes:

In the case of ontological disputes, when one distinguishes senses for the problematic terms, the disputes appear to persist as strongly as ever. For example, even once one distinguishes ‘exists_n’ and ‘exists_u’, the nihilist and the universalist will continue to disagree, for example over questions such as ‘Are there any u-objects that are not n-objects?’ And even once one distinguishes ‘exists_n’ from ‘exists_u’, the nihilist and the universalist will continue to disagree over questions such as ‘If there exists_u an X, does an X really exist?’ from ‘There are things arranged X- wise,’ the nihilist and the universalist will continue to disagree, for example over questions such as ‘If there are things arranged X-wise, does an X *really* exist?’

(Chalmers 2009, p. 89-90)

Hence, it certainly remains an open question whether, provided that one could somehow ensure that all the ontologists spoke the exact same language, they would suddenly unanimously agree on all the ontological facts. And even if they would agree on all the ontological facts, we can still question whether these facts are somehow without substance or trivial. If ontological disputes turn out to be merely verbal, this would therefore at best show that these specific disputes are misguided or shallow; not that ontological inquiry as a whole is trivial or without substance.

1.2 Thomasson’s *Easy Ontology*

1.2.1 *Easy Ontology and Metalinguistic Negotiations*

A metaontological view which in some respects can be seen as an improvement on Hirsch’s Verbal Disputes View is Amie Thomasson’s *Easy Ontology* (2007; 2009; 2014; 2017). The view defended by Thomasson is roughly a contemporary take on Carnap’s external/internal distinction, which was mentioned in the introduction. On the semantic level, Thomasson argues that ‘exists’ has a fixed, formal rule of use which is intended to correspond to Carnap’s internal use of the quantifier. She calls this ‘schema E’ (Thomasson 2014, p. 86):

E : Ks exist iff the application conditions actually associated with ‘K’ are fulfilled.

(As far as I can tell, Thomasson conceives of application conditions roughly in the same fashion as I conceive of rules of application.) When ontological questions are asked in the internal sense, this has as a consequence that ontological questions are generally too easy to answer to warrant serious metaphysical debate. Whether the application conditions for any given term ‘K’ are fulfilled is generally easily settled by means of empirical inquiry or conceptual analysis. As a result then, ontology is “easy” -- at least when we stick to asking internal ontological questions. Philosophers, however, tend to abuse this formal rule of use, and attempt to ask external-ontological questions. Hence, when the philosopher asks whether tables exist, they intend this question not to be a trivial question about whether there is anything which satisfies the application conditions for ‘table’ (as specified by our linguistic framework), but rather whether employing the term ‘table’ somehow adequately reflects how things are. The problem is that this external-ontological question is guilty of a use-mention fallacy, and therefore confused. To elaborate, when the question is asked whether tables exist in the internal sense, the term ‘table’ is legitimately *used*. Ontologists, in contrast, are attempting to ask a question which merely *mentions* the term table, i.e. whether ‘tables’ exist, which makes no sense if construed as a question about the world. The only way to make sense of this question is as a pragmatic question about language, i.e. whether ‘tables’ *should* ‘exist’, in other words, whether we should keep using the term ‘table’, or whether ‘table’ is a useful concept. According to Thomasson, then, philosophical ontology is largely a confused project, which, at least in its current form, ought to be dismissed (Thomasson 2014).

Although it does not rely on quantifier variance, Thomasson’s view faces similar objections to the verbal disputes view. Because it portrays serious ontologists as attempting to ask confused pseudo-questions, it ostensibly fails to provide a charitable account of what these serious ontologists are up to. If ontological questions are so easily answered, why have ontologists been so confused for so long, and why are these ontological debates still ongoing? Just like the verbal disputes view, the easy approach to ontology appears highly uncharitable to what some of the serious ontologists are doing. More specifically, it appears that because ontological truths are so easy to come by, anyone doing ontology must be asserting trivialities. Although this does not pose much a problem for explaining what the commonsense ontologists are up to -- even if they’re asserting trivialities, somebody obviously still disagrees -- the view comes out as highly uncharitable to any “revisionist” ontologists, as it portrays them as asserting trivial falsehoods. Easy ontology therefore also needs to offer a plausible account of what the

ontologists which are apparently asserting trivial falsehoods, such as ‘tables don’t exist’, could possibly be hoping to achieve. As Thomasson herself puts it:

[W]e can make some sense of why one would utter a trivial truth in contexts in which it has been publicly denied—so the one who utters a trivial truth may often be seen as doing something sensible at least in the context of such debates. The tougher side of the challenge is making sense of what disputants who utter a trivial falsehood can be doing. They can’t be treated as incompetent speakers (at least, they seem to use terms in the normal way outside the philosophy room) or as making perceptual errors. And so it does seem that the easy ontologist faces similar difficulties in charitably interpreting what such serious ontologists are up to—not in interpreting what they are saying but what they are doing.

(Thomasson 2017, p. 9).

In response to this worry, Thomasson has recently argued that we can improve on the verbal disputes view by appealing to the idea of *metalinguistic negotiations* (Thomasson 2017). The notion of a metalinguistic negotiation was first introduced by Plunkett & Sundell (2013) and refers to a normative metalinguistic dispute, i.e. a dispute in which the disputants use a term instead of mentioning it in order to express a view about how the term in question should be used. In other words, a metalinguistic negotiation is a dispute about the meaning of a term in which the participants are conducting their dispute at the level of pragmatics rather than at the level of semantics; by using the term in their preferred way, they express a view about how the term *ought* to be used.

For example, imagine a dispute about whether waterboarding is torture (Chalmers 2011; Plunkett & Sundell 2013). Participant A says ‘Waterboarding is torture’ and participant B says ‘No, waterboarding is not torture’, and participant A and B are using different definitions of ‘torture’. Let’s assume that participant A is adhering to something roughly along the lines of the United Nations’ definition of torture, on which any act inflicting severe suffering for the purposes of extracting information or as punishment is torture, and participant B endorses something similar to the United States Justice department’s former definition of torture as any act of inflicting severe pain resulting in death, organ failure, or permanent damage of a significant bodily function. Considering their different definitions, participant A and B thus mean something different with ‘torture’ and their claims about waterboarding aren’t actually in conflict with each other. According to Hirsch’s definition, A and B would in fact appear to be having a verbal dispute. However, although participant A and B do not disagree with each other on the semantic level, they might still genuinely disagree with each other about how the word ‘torture’ ought to be used. Hence, participant A might keep insisting that “Waterboarding is torture” as

way of insisting that we ought to adopt a more inclusive definition of torture, to ensure that abhorrent practices such as waterboarding remain outlawed on the same grounds as other cruel acts carried out for the purposes of punishment or obtaining information. Similarly, participant B might keep insisting that “Waterboarding is not torture” as a means of advocating for a narrower definition of torture, perhaps because they believe that a narrower definition will help steer more governments away from practices that induce lasting physical harm. It is as such possible that there is a genuine dispute taking place between A and B, although there is no disagreement between them on the semantic level.

Thomasson (2017) argues that we can make sense of what is going on in some metaphysical disputes if we understand them as metalinguistic negotiations of this kind. (In principle, this appeal to metalinguistic negotiations is also open to the proponent of quantifier variance (see Belleri 2017)). For example, when metaphysicians argue about the persistence conditions of works of art, or of persons, they are arguing about which definitions of ‘art’ or of ‘person’ that we ought to adopt, seeing as these definitions have normative consequences for e.g. how we classify and evaluate works of art, or for how we treat human fetuses or human individuals in persistent vegetative states. This understanding of metaphysical disputes as metalinguistic negotiations thus lines up neatly with Carnap’s conception of external questions as pragmatic questions about which linguistic framework to adopt. In the case of ontology, Thomasson argues that some traditional debates about existence can be seen as “reflecting disagreements not about *how* the term should be used, but rather about whether the term *should be used at all*” (Thomasson 2017, p. 16, emphasis in the original). So for instance in debates about whether there is anything such as free will, or about whether race exists, the eliminativists can be interpreted as advancing the view that we stop using the concept of free will, or the concept of race, because these are pragmatically harmful concepts. Similarly, in the debate between the universalists and nihilists about the existence of ordinary objects, the nihilists could according to Thomasson be viewed as advocating for the pragmatic stance that we should stop using concepts such as ‘table’ or ‘book’ because they do not function properly, as their use involves us in (alleged) contradictions. As an example, Thomasson refers to Trenton Merricks’ argument against the existence of ordinary objects, which states that because tables and books must be causally efficacious in order to exist, they cannot exist because if they did, they would be causally redundant. So according to Merricks, assuming that tables and books exist therefore involves us in contradiction (Merricks 2001).

On Thomasson's view, then, the metalinguistic negotiations account is able to provide a more charitable interpretation of what the serious ontologists are up to, as metalinguistic negotiations are in many contexts certainly worth engaging in. Because many of our concepts are bound up with normative issues, the concepts that we employ can have far reaching practical implications. The metalinguistic negotiations view of ontological disputes can also do more justice to the serious ontologists' intuitions that their disputes are not "just about language" but about the real world. Metalinguistic negotiations about how we ought to employ words such as 'person', 'race', or 'torture' are very much world-focused, as they have practical societal and institutional implications. At the same time, it allows us to maintain the view that the ontologists are genuinely disagreeing and not only talking past each other, as there is a real dispute going on at the metalinguistic level. The metalinguistic negotiations account, then, is clearly an improvement over the verbal disputes view for making sense of ontological debates.

1.2.2 Problems for the Metalinguistic Negotiations Account

There are nonetheless some serious problems for the interpretation of ontological disputes as metalinguistic negotiations. Although I agree with Thomasson that the metalinguistic negotiations account does a good job of making sense of metaphysical debates about 'race' or 'art', and to some extent 'person' (more on this below) -- debates typically thought of as concerned with *social* ontology -- the prospects for understanding more traditional ontological debates about the existence of numbers, time, or universals are much less promising.

Firstly, it should be observed that contrary to Thomasson's proposal, the metalinguistic negotiations account as a metaontological view isn't actually compatible with her object-level account according to which 'exists' has a formal, invariant rule of use. Thomasson argues that one of the main appeals of easy ontology is that we can make sense of ontological disputes as non-verbal disputes in which the meaning of 'exists' remains fixed (Thomasson 2015; 2017). However, the metalinguistic negotiations account requires for it to work that the disputants give different meanings to the term that is being negotiated -- otherwise there's simply nothing to negotiate. Hence, for a metalinguistic negotiation to occur over whether "tables exist" or whether "numbers exist", the disputants must be able to assign different meanings to these utterances. It therefore seems that the metalinguistic negotiations account of ontological debates in fact needs a notion of quantifier variance in order to work, or it needs to insist that it is the terms for the things that are said to either exist or not exist that are being negotiated, such as

‘number,’ ‘race,’ ‘person,’ ‘universal’ or ‘table.’ While it seems plausible that in disputes over whether race exists, the participants are in fact using different definitions of ‘race’, it sounds quite implausible that the serious ontologists arguing about the existence of numbers or universals or ordinary objects give different meanings to the words ‘table’, ‘number’ or ‘universal’. For this reason, in order to make sense of the more traditional ontological debates as metalinguistic negotiations, it seems that one would need to supplement this view with quantifier variance, which is in direct conflict with Thomasson’s first-order conceptual analysis of existence.

Naturally, Thomasson could resist this move and argue that her schema E describes the correct use of the concept of existence in ordinary English. This is compatible with the view that the serious ontologists are *attempting* to use ‘existence’ to mean something different -- they are simply in error, misusing their words. However, it would then no longer constitute an improvement over the verbal disputes view as a metaontological account, as it would then succumb to similar objections to the verbal disputes view and more. In particular, it would be stranded with an implausibly uncharitable account of what the serious ontologists are up to, as they must all be uttering obvious falsehoods or nonsense (not even truths in their own language!). Further, easy ontology would be at great pains to explain why, if our ordinary use of the existence-concept is the only use which we can truly make sense of, and it is non-defective, these philosophers feel the need to initiate metalinguistic negotiations? Why think that we ought to use the concept of existence differently, if there is nothing at fault with how it currently functions?

It might then seem preferable to accept quantifier variantism and try to improve on the verbal disputes view by also appealing to metalinguistic negotiations. Easy ontology would then just amount to an improved version of the verbal disputes view with an added analysis of what the correct use of ‘exists’ in ordinary English. The notion of a metalinguistic negotiation could then also aid the verbal disputes view in countering some of its objections. For example, it would provide an explanation for why ontological disputes are so persistent, as they could then in principle be characterised as reflecting deep normative disagreements about how ‘exists’ ought to be used.

The resultant view does unfortunately run into some new problems of its own. The main difficulty is in explaining exactly what the normative significance of how we use the concept of existence would be. In the case of the meaning of ‘race’, or ‘art’, or even ‘free will’, the normative implications are relatively straightforward, but what could the deep normative disagreements be

which underlie the metalinguistic negotiation about existence or quantificational language? Why would anyone engage in such a metalinguistic negotiation?

Let's take as an example the ontological debate about composition, or about the existence of ordinary objects such as tables. In these debates, the serious ontologists who argue for nihilism about composition or eliminativism about causally redundant entities still admit of "particles arranged tablewise" or that, for all practical intents and purposes, it is better to keep talking of "tables" as if they existed, or as if other composites corresponding to middle sized objects really did exist. There thus seems to be absolutely no pragmatic or normative import at all of arguing that *strictly speaking*, "tables" and other composites do not exist. Rather, the concern of the serious ontologists appears to be with what description truly reflects or better represents reality, the way things *really are*, independently of any practical concerns. At least, that appears to be the self-conception of most serious ontologists, and it seems to be the only way of charitably explaining why they are so invested in these disputes.

Now, the advocate of the metalinguistic-negotiations-as-metaontology account can of course grant that this self-conception is correct, but that this is a confused enterprise; all linguistic frameworks/conceptual schemes are descriptively equally good -- there is no metaphysically or epistemically privileged "best" language for describing or conceptualising the world. And so we conclude, with Carnap, that the only sensible philosophical ontological questions we can ask are pragmatic ones.

However, it seems to me that such an inference from the claim that all conceptual schemes are equally good to the statement that there are no non-pragmatic ontological questions does not follow. To clarify why that is so, let me turn to an example.

Take as an example the metaphysical debate over personal identity; this is one of the debates which Thomasson (2017) suggests qualifies as a metalinguistic negotiation. Imagine the following thought experiment from the personal identity literature, in which at some point in the future, an advanced human civilisation have invented a teletransportation device. The teletransporter works as follows:

[W]hen you press a button, a scanner records the states of all of the cells in your brain and body, destroying both while doing so. This information is then transmitted at the speed of light to some other planet, where a replicator produces a perfect organic copy of you. Since the brain of your Replica is exactly like yours, it will seem to remember living your life up to the moment when you pressed the button, its character will be just like yours, and it will be in every other way psychologically continuous with you.

(Parfit 1987, p. 21).

Now imagine that, despite their technological advancedness, the philosophers and scientists of this human civilisation have not yet reached any consensus on what a “person” is. Due to the recent invention of teletransportation, they are however in urgent practical need of an adequate conception of personal identity which they can employ in laws and regulations regarding teletransportation. After a period of intense public debate, the people in charge of government advised by an expert commission decide to settle on the conclusion that the replica is the same person as the person who stepped into the teletransporter (provided that there is only one replica, making multiple replicas is strictly forbidden and punishable by law), and the replica will have just the same legal status, rights, properties, assets and so on. For all practical intents and purposes, the replica is the same person; if somebody pre-transportation incurred a great debt, a prison sentence, or is the legal parent of a child, the replica will post-transportation be the owner of the debt, the bearer of the prison sentence and the parent of the child. According to the metalinguistic negotiations account, this sort of pragmatic decision making about when we should take somebody to be “the same person” is all there is to the debate about personal identity.

However, we could imagine that there is a group of philosophers who refrain from ever using teletransportation, despite the fact that there are no legal complications, because they are convinced that the person coming out on the other side will not be *them*. These philosophers think that, effectively, entering the teletransporter amounts to a kind of death; the person on the other side will simply be somebody else that happens to be molecule-for-molecule exactly similar. Now, these philosophers might not desire a change in laws per se, as they take no issue with other people who want to use the teletransporter doing so. They perhaps think of it as a form of suicide which one could easily commit without incurring any significant cost to society or one’s loved ones (the replica will presumably be just as suicidal, but at least now that’s *their* problem, not yours) -- although legally speaking, you are of course still alive. Thus, for all practical intents and purposes, they believe that the replica may be treated as the same person, although “strictly speaking,” it will not be.

Now, the concern of these philosophers is presumably not normative or pragmatic at all. When they insist that “the replica is not the same person”, they do not necessarily do so because they believe that we ought to use the concept of a “person” differently -- at least not for any

practical reasons -- but they do so on purely descriptive grounds. When these philosophers disagree with the law which says ‘the replica is the same person’, they disagree because they think that saying that it is the same person poorly reflects some natural facts, not because the law needs to be changed. It would therefore be inaccurate to describe these philosophers as engaging in a metalinguistic dispute about how we ought to use the term ‘person’ -- rather it seems that they are attempting to make a substantive claim about the nature of reality.

To make explicit that this is certainly not a dispute about language, we can even imagine that these dissenting philosophers concede the term ‘person’ to the lawmakers. For example, we can make a distinction between *legal personal identity* and *subjective personal identity*. Legal personal identity refers to the persistence conditions of persons as set out by the law, and subjective personal identity corresponds to that which the group of philosophers believe dies (in a non-legal sense) when one enters the teletransporter. They can then agree to use ‘person’ to refer to legal personal identity, and continue on in their inquiry about the persistence conditions for subjective personal identity.

What this example illustrates is that although normative concerns are certainly important components in many ontological debates, it is not plausible to think that *all* ontological debates are exhausted by these practical concerns. Of course, one might wish to dismiss the dissenting philosophers’ inquiry as confused -- perhaps because one thinks that there is no such thing as “subjective personal identity” -- but also this dismissal amounts to a substantive ontological thesis. Thus, whether subjective personal identity exists or not, the question *whether* it exists is neither a pragmatic nor an ontological pseudo-question. The dissenting philosophers’ inquiry into the persistence conditions (if any) of subjective personal identity is a substantive and non-trivial ontological issue which does not turn on exclusively practical concerns.⁸

Moreover, appealing to the thesis that “all conceptual frameworks are equally good” does not help to deflate the ontological question about subjective personal identity. The claim is not that the dissenting philosophers’ concept of subjective personal identity is somehow *better* than the concept of legal personal identity; they are simply different concepts of personal identity. In fact, each concept fills a rather different function, and which concept we employ depends on our

⁸ One might object that the question of whether subjective personal identity exists is, on Carnap’s distinction, in fact an internal question. However, this move does not help the proponent of the external/internal distinction as a means of deflating ontological debates, as clearly the question about subjective personal identity is highly non-trivial and not easily settled, despite being an internal question. In that case, the example would instead refute the claim that all internal ontological questions are easily answered by conceptual or empirical means, again supporting a non-deflationary position.

practical or epistemic interests. The deeper point is that in many cases we can just introduce a new concept for that which we are interested in talking about. Asking whether one concept is better than the other is a little bit like asking whether the concept of orange is better than the concept of red; there's no comparison to be made, they are simply different concepts. Which concept is better depends on whether we are interested in making claims about things that are orange or things that are red. As such, there is no "better" simpliciter; the only sensible question to be asked here is "Better for what?"

I think we should think of quantifier variance in the same way. Different existence-like concepts, or different senses of the quantifier, should simply be thought of as different concepts. These different concepts pick out different things, and which one we employ will depend on our interests. None of these concepts (provided they are all non-defective, of course) are necessarily any better than any other; they just have different pragmatic and epistemic roles. Thus, before we can ask our non-pragmatic questions, we need to make a pragmatic choice about which concept to use. However, the fact that which concept(s) we ultimately choose to use matters only pragmatically does not mean that ontological questions are shallow or exclusively pragmatic questions; it means that the nihilists and the universalists, in so far as they are using different concepts, are asking different questions. Plausibly because they are interested in different things... For these reasons, an appeal to the fact that different linguistic frameworks or conceptual schemes are equally good does very little to establish that ontologists are asking pseudo-questions.

In summary, then, easy ontology does not provide an adequate interpretation of what serious ontologists are up to. Firstly, in order for the metalinguistic negotiations account to constitute an improvement over the verbal disputes view of traditional ontological debates, it needs to rely on quantifier variance. Secondly, even if invoking metalinguistic negotiations can help us explain some (social) ontological debates over for example race, gender and artworks, the more traditional ontological debates are still unaccounted for. Finally, appealing to "equally good" linguistic frameworks or conceptual schemes does not serve to deflate ontological debates, as it fails to show that the serious ontologists are asking pseudo-questions.

1.3 Chalmers' Ontological Anti-realism

1.3.1 Ontological Anti-realism

As we have just seen, neither Hirsch's verbal disputes view nor Thomasson's easy ontology provide an adequate explanation of what is going on in ontological debates. Although it appears plausible that some ontological debates could turn out to be merely verbal, or metalinguistic negotiations, this certainly does not apply to all debates, and more significantly, it does not warrant a deflationary stance towards ontology in general.

David Chalmers' has defended yet another species of ontological deflationism; *ontological anti-realism*. Ontological anti-realism is the view that ontological assertions lack determinate and objective truth-values. An assertion has a determinate truth-value when it is either true or false, and its truth-value is objective when it is the same in all contexts of assessment (Chalmers 2009). Hence, ontological anti-realism amounts to the view that ontological claims, at least sometimes, are neither true nor false, or can come out as either true or false depending on the context in which the claim is assessed. On an anti-realist view, then, a claim such as 'numbers exist' might be indeterminate, or true in some contexts but not in others (e.g. it might be true in a mathematician's office but false in the philosophy room).

We can also make a distinction between weak forms of anti-realism and stronger forms of the same view, depending on the proportion of ontological claims which come out as indeterminate or context-sensitive. A *weak* ontological anti-realism would hold that there are at least some unproblematic ontological existence assertions that lack objective and determinate truth-values, while the absolute strongest form of ontological anti-realism would hold that *all* unproblematic ontological existence assertions lack determinate and objective truth-values. Then there is also room for all kinds of moderate anti-realisms in between, and one could even be an anti-realist about a specific class of ontological existence assertions (e.g. assertions about ordinary objects, or numbers). Plausibly though, a fairly strong form of anti-realism could still allow for a limited class of ontological existence assertions that have determinate and objective truth-values, such as 'there exists square circles' or 'there are concrete unicorns', which appear to be false in all contexts of assessment quite independently of any specifically ontological considerations (Chalmers 2009). Chalmers himself appears to favour a fairly strong anti-realism of this latter kind.

Moreover, Chalmers defends a type of contextualism about quantification. According to Chalmers, there are several different contents which our concept of existence can express, and which one is expressed is determined by the context. When we make ordinary existence claims in everyday contexts, like 'there is an apple on the counter', we use the quantifier in a "lightweight"

sense, and our assertions are determinately either true or false depending on which content is expressed. The truth-values of such ordinary existence claims are generally trivially established by empirical or conceptual means.

Contrastingly, Chalmers argues, when we make existence claims in the ontology room, we are using the quantifier in the “heavyweight” sense which attempts to express a concept of absolute quantification. This so called *absolute existential quantifier* is supposed to be independent of any particular context. Unfortunately, the absolute existential quantifier does not have a determinate content, which renders it indeterminate which content, if any, is actually expressed. Questions and assertions intended to express the absolute existential quantifier are therefore all defective utterances. Because it is indeterminate what the contents of our quantificational expressions inside the ontology room are, these ontological claims fail to have determinate truth conditions, making it the case that ontological claims are neither objectively nor determinately true or false. Chalmers assimilates these defective heavyweight uses of the quantifier to Carnap’s external questions, while the non-defective lightweight uses are comparable to Carnap’s internal questions.

In order to further precisify his view, Chalmers introduces some technical apparatus. The way that quantificational claims work, according to Chalmers, is that they are made at a *world* with respect to a *domain*. The world, however, does not come with a pre-specified domain; the domain that the quantificational claim should be evaluated with respect to is determined by its context of utterance. This mapping from worlds to domains he calls a *furnishing function*.

A furnishing function (or equivalently, a *domain-determination function*) is a mapping from worlds to domains. A world and a furnishing function jointly determine a furnished world. In effect, given a world, the furnishing function specifies a class of entities that are taken to exist in that world.

(Chalmers 2009, p. 108)

The context of utterance thus determines a furnishing function, which together with the world determines the domain. If the thing quantified over is part of the determined domain, then the thing in question exists. The domain is as such what enables us to assign truth-values to quantificational claims.

In this terminology, what the serious ontologists are attempting to do when they use the absolute existential quantifier, is to quantify over an absolute domain; the domain of everything that exists. The anti-realist, contrastingly, maintains that there are several admissible domains or furnishings. When the serious ontologists attempt to ask questions using the absolute existential

quantifier, they attempt to do so independently of any particular furnishing, which yields an indeterminate domain. The realist who holds on to the view that existence claims using the absolute existential quantifier must therefore deny that there are several admissible domains, and maintain that the world only has one admissible domain or furnishing (i.e. the absolute domain).

The only method for evaluating ontological claims made inside the ontology room is to assign truth values by *supervaluating* on all admissible furnishing functions. Hence, only those ontological assertions which would come out as always true or always false on all admissible furnishings, are determinately true or false (similarly to a method for assigning truth-values to vague claims). For example, ‘There are square circles’ would plausibly come out as false on all admissible furnishings, whereas ‘Numbers exist’ would come out as true given some furnishings and false given others, meaning that the former is determinately false while the latter has an indeterminate truth-value.

As a consequence, Chalmers endorses a type of quantifier variantism or ontological pluralism about lightweight quantification, as the content expressed by ‘exists’ varies between ordinary contexts. Nonetheless, Chalmers is still a quantifier monist about heavyweight quantification, as in the ontology room ontologists intend to express a concept of absolute existential quantification when they use the quantifier. The reason why these heavyweight ontological claims generally fail to have determinate truth values is due to the fact that the concept of absolute existential quantification appears to be defective. He writes:

I do not think it is obvious that there *is* a concept of absolute existential quantification. What I think is obvious is that ontological discourse functions as if there is such a concept. It is plausible that many ontologists are tacitly *intending* to use such a concept when they make ontological existence assertions, and indeed when they think ontological thoughts. It may be that there is something defective about these attempts, in which case there may only be a *meta-concept* in the vicinity: that is, a concept of the concept of absolute existential quantification. We might also say that there is at least a *pseudo-concept* in the vicinity: something that functions in our thought and talk like a concept, in some respects, while falling short in other respects (including respects tied to truth-evaluability, perhaps).

(Chalmers 2009, p. 91-92)

The main virtue of Chalmers’ anti-realist picture is thus that it can retain the idea that serious ontologists are attempting to make substantive, non-trivial claims about the world. This aligns much better with the serious ontologists’ self-perception, as it doesn’t require that ontologists are making claims that are trivial, merely talking past each other (and so not really disagreeing), or that they are only disagreeing about how the quantifier ought to be used. Instead we can allow

that the serious ontologists are all using, or attempting to use, the very same concept (i.e. that of absolute existential quantification) -- even though there is no non-defective, first-order concept of this kind. But at least all the serious ontologists might share a non-defective meta-concept of the concept of absolute existential quantification, which explains why it appears to them that they are arguing about the very same thing. Chalmers' view thus appears to be a significantly more charitable way of deflating ontological debates.

It is, nonetheless, still a deflationary account of ontological debates, as on Chalmers' view, the serious ontologists are attempting to answer questions that simply have no determinate or objective answers. In a sense, these are pseudo-questions, as it isn't determinate what the contents of the questions themselves are. In effect, Chalmers' anti-realism comes out as a type of error-theory, as the serious ontologists erroneously take there to be objective and determinate questions about what exists in some absolute sense when there are no such questions. The lack of objective and determinate utterances involving the absolute existential quantifier has the result that ontological debates are entirely pointless, as it is indeterminate what these debates are actually about.

1.3.2 Objections to Anti-realism

Although Chalmers's anti-realist account could be seen as an improvement over the verbal disputes view or metalinguistic negotiations account, it is not without its problems. Firstly, there is some unclarity in the notion that the absolute existential quantifier is defective, or merely a pseudo- or meta-concept. Secondly, it seems implausible that there could genuinely fail to be ontological facts, and as a matter of fact, this anti-realist thesis does not seem to follow from Chalmers' semantic analysis of the absolute existential quantifier as a defective concept. I will discuss these objections in turn.

Firstly, one might be curious about exactly what it is that makes the absolute existential quantifier defective. Evidently, part of its defectiveness is that its meaning is somehow indeterminate, but clearly it must have something that at least *resembles* a determinate meaning -- given that nobody would take it to express for instance the concept of swimming, or of womanhood. Chalmers suggests that there is perhaps a pseudo- or meta-concept, but he leaves it up to the reader to create their own idea of what these are. He further remarks that "[i]t is tempting to hold that the absolute quantifier is something of a philosopher's invention, one that otherwise plays very little role in our thought and talk" (p. 103). But even if it is a philosophers'

invention, this does not automatically explain why its meaning should be defective or indeterminate; scientists introduce new scientific concepts into their theories all the time that seem to be perfectly functional and determinate.

An additional issue is explaining why, if the meaning of the absolute existential quantifier is defective, the serious ontologists do not recognise this, but keep using it as if it had a determinate meaning. Plausibly, it cannot be indeterminate to the disputants what *they* mean when they use the (pseudo-)concept, as otherwise it seems difficult to explain how they come to the conclusions that they do, and why they have confidence in their claims. Further, if there are no objective or determinate ontological facts, does that make ontology purely a matter of aesthetic preference, or personal taste? Considering my conclusion from the previous section -- that traditional ontological results have very little, if any, normative significance -- what could be at stake for these philosophers?

More importantly, not only are there some loose ends to tidy up assuming that the concept of absolute existential quantification is defective; the main problem with Chalmers view is that the anti-realist conclusion that he draws from the first-order analysis of the absolute existential quantifier as defective does not follow. Chalmers wants to conclude that there are no objective or determinate facts of the matter about existence. But if *that's* the conclusion that Chalmers wants to draw, it is not clear where the indeterminacy is supposed to come in. Chalmers argues that, due to a defect in the concept of absolute existential quantification, utterances involving this concept fail to have a determinate content. This makes it sound as if the indeterminacy is supposed to be strictly a matter of language (as opposed to the worldly ontological facts somehow being indeterminate). The problem is that if the indeterminacy is merely a property of the ontologists' quantificational language, then it is not obvious why the anti-realist conclusion should follow.

In particular, it is highly unclear why the fact that ontologists fail to properly express their questions and theories using precise and determinate concepts should imply that no determinate or objective answers are available if the questions were to be made precise and determinate. For example, we can imagine somebody trying to describe a specific shade of colour over the phone, while for instance in the process of buying some clothing. Unfortunately, because this particular person's colour vocabulary is limited (perhaps they are not a native speaker of the language), they aren't doing a very good job of describing the specific shade of colour, and the person they are trying to communicate it to isn't sure which exact shade they are

describing. This, however, doesn't imply that the shade of colour is *indeterminate* -- the person doing the describing might have a perfectly determinate shade in front of them. It is merely indeterminate to the person on the other end *what it is that they mean*. In most contexts, failures of expression, communication, or interpretation are usually seen as just that; not as symptomatic of some kind of indeterminacy of the subject matter. The way to resolve these failures of communication is not to give up and conclude that there is no fact of the matter, but rather to communicate better, or introduce more precise concepts.

Therefore, it seems that in order to establish the anti-realist conclusion, the indeterminacy must be somehow located in the world and not just in language; the facts themselves must on Chalmers' view be indeterminate. But why should we believe, on the basis of an observation that we are using a defective concept, that the facts are indeterminate? It appears that additional argumentation would be needed to support or clarify the anti-realist position.

A possible response on behalf of the anti-realist could be that heavyweight quantification is so defective that there are just no wordly facts in the vicinity which we could plausibly take as candidates for ontological discourse to be about. Hence, we do not need to make the claim that the world is indeterminate, as absolute existential quantification is not about the world at all. But given that lightweight quantification is non-defective, and by Chalmers own lights unproblematically succeeds in making determinate claims about the world, this is not a very persuasive reply. Evidently there are determinate facts in the vicinity of the facts that the serious ontologists are attempting to express when they use heavyweight quantification, even if it is indeterminate which precise facts they are intending to express.

A third possible option for the anti-realist would be to locate the indeterminacy in the epistemic realm. This could be accomplished by saying that even if there are determinate and objective facts about what exists in some absolute sense, we can never know these facts. Therefore, existence facts are *epistemically* indeterminate. Such a view, though, would strictly speaking be *epistemicism* (see Bennett 2009 for a development of a view along these lines), not anti-realism. Epistemicism is entirely compatible with ontological realism, and potentially incompatible with anti-realism.

In Chalmers' more technical terminology, we can put the point as follows. Chalmers argues that the absolute existential quantifier fails to pick out a unique or determinate furnishing function. As Chalmers himself recognises, the technical apparatus does not help us choose between realism or anti-realism, as whether there are several admissible furnishing functions or

merely one (i.e. whether the world comes with a built-in domain), is a substantive ontological question which must be independently settled; that there is more than one admissible furnishing function is a mere stipulation. However, my point is that, even if we grant that there are several admissible furnishing functions, this need not lead to anti-realism. On Chalmers' picture, the different furnishing functions are meant to correspond to the contextual variations in meaning of the (lightweight) existential quantifier. That is, the quantifier has different meanings in different contexts, and in the ontology-context it is indeterminate which of all its possible meanings it has. This means that the view Chalmers is advancing could be construed as a form of quantifier variantism, as he himself acknowledges (Chalmers 2009, p. 123). But if Chalmers is relying on quantifier variantism, Chalmers appears to be guilty of the same fallacy as the critic of quantifier variantism who takes quantifier variantism to imply anti-realism or relativism (as I argued in section 2.1.2). As I acknowledged previously, there is of course a trivial sense in which existence claims are relative to which language we speak, i.e. to which existence-like expression we are using (and similarly, on Chalmers' view, the content of existence claims are relative to a furnishing function). That being said, this is a very shallow type of *linguistic* relativism, and to my mind, it would be misleading to call this *ontological* relativism or anti-realism.

Let's go through this in more detail: if the lightweight quantifier has different meanings in different contexts, then this means that a sentence such as 'tables exist' is going to express different contents in different contexts. In fact, by uttering the sentence 'tables exist', the claim that is made will vary between contexts. Hence, in context A we might be making the claim that "tables exist_a", and in context B that "tables exist_b". Nonetheless, the truth-values of the propositions 'tables exist_a' and 'tables exist_b', will still be fixed by facts in the world, which obtain independently of the context (unless, of course, we postulate some kind of indeterminacy in the worldly facts, which we already established that it seems that Chalmers isn't trying to do). But in that case, it is not accurate to say that the truth-value of the proposition that "numbers exist" is relative to the context. What *is* relative to the context is which proposition is expressed by the sentence 'numbers exist', because we can use the same sentence to express different propositions in different contexts. As such, this is a type of linguistic relativism rather than ontological relativism or anti-realism. It is the meanings of our words that are relative to the furnishing function; not the facts (the facts, I take it, are independent of which language we speak).

Now, of course, Chalmers wants to maintain that the meaning of the *absolute* existential quantifier does not vary in meaning like ordinary "lightweight" quantifiers. The ontological

context is one in which philosophers are attempting to use the quantifier in some absolute, context-independent sense, which makes it indeterminate what the absolute existential quantifier means because it does not pick out a determinate furnishing function. Yet, the same point holds: it does not follow that the ontological facts are therefore indeterminate. What is indeterminate is merely exactly which of all possible propositions is expressed when we utter the sentence ‘tables exist’ inside the ontology room. If we were to make it precise exactly which content is expressed, it is highly unclear how the claim could fail to have a determinate, context-independent truth-value.

The upshot of Chalmers’ observations should as such be that we need to always be specific, also inside the ontology room (or *especially* inside the ontology room), about with respect to which furnishing function we are making our ontological claims. If there are several admissible furnishing functions, that does not mean that ontological facts are indeterminate, it just means that there are many different ontological questions that can be asked and answered, and that we need to be more precise about which one it is we are asking. Anti-realism thus seems to me to be the wrong conclusion here; if the absolute existential quantifier fails to pick out a determinate content/furnishing function, this is a failure of language, not a symptom of there being no objective or determinate ontological facts.

In conclusion then, Chalmers semantic analysis of heavyweight quantification as defective does not warrant any anti-realist conclusions. More specifically, it is highly unclear why our failure to precisely express ontological claims in language such that they have determinate meanings should imply that there are no objective or determinate facts of the matter. In fact, even if ontological discourse as it plays out in the ontology room turns out to be defective just in the way that Chalmers proposes, this is entirely compatible with ontological realism.

1.4 Conclusion

In this chapter I have argued that none of the three deflationary metaontological views that I have discussed offer a satisfactory metaontological account. Both the verbal disputes view and the metalinguistic negotiations account fall short in their interpretation of what the serious ontologists are up to on the basis of considerations of charity. Chalmers’ anti-realism does better in this regard, but also raises some questions. Most significantly, none of the views succeed in deflating ontological debates, as none of the views succeed in showing that there are no

objective, determinate or non-trivial ontological truths to uncover. Even if these deflationary accounts may correctly diagnose some current ontological debates as merely verbal, linguistic negotiations or otherwise linguistically defective, *in principle*, the possibility of engaging in substantive and non-defective ontological inquiry and debate stands unscathed.

2.

INCONSISTENT CONCEPTS

This chapter is a detour into the theory of inconsistent concepts. I hope to make salient what the relevant phenomenon is and present a preliminary framework for thinking about inconsistent concepts. Further, I aim to lend plausibility to the idea that there really are such things as inconsistent concepts, and respond to some recent objections to this idea. The discussion will proceed as follows. Firstly, in section 2.1, I briefly discuss what concepts are hypothesised to be in the recent literature, and where the discussion of inconsistent concepts might fit into all this. In 2.2 I offer a more detailed account of what inconsistent concepts are, and how inconsistency differs from other conceptual defects. Section 2.3 addresses concept possession, and the potential problems concept possession presents for the theory of inconsistent concepts. In addition, I argue for my preferred account of concept possession. Finally, 2.4 responds to some objections from error theory and externalism to the idea that there are inconsistent concepts.

2.1 What Are Concepts?

Concepts are controversial entities, with multiple possible functions in thought and language. Generally, there are three distinct roles which concepts are called upon to fill:

Firstly, concepts are commonly hypothesised to be the constituents of thought. In turn, if thought has a propositional structure, concepts are also the basic constituents of propositions. On this perspective, concepts are roughly the mental analogues of words; just as we combine words in natural language to make up sentences, we combine concepts in our minds to make up thoughts. The role of concepts as the constituents of thought has most notably been emphasised and theorised by Jerry Fodor (e.g. 1975; 1998). Additionally, it is a tacit assumption in much contemporary philosophy of mind that this is the primary role of concepts.

Secondly, concepts are often appealed to as the meanings of words, or more specifically the aspect of meaning which is not exhausted by reference. A natural way to elaborate on this

idea is to say that concepts are Fregean senses. To illustrate, the names ‘Hesperus’ and ‘Phosphorus’ both happen to refer to Venus, but they express different senses or concepts, namely those of HESPERUS and PHOSPHORUS respectively. Another way of putting this idea is to say that concepts are intensions, as opposed to extensions. For instance, the words ‘cordate’ and ‘renate’ happen to have the same extensions, because in the actual world every creature with a heart also has a kidney (at least if philosophers are to be trusted on these matters), but they clearly express different concepts; i.e. that of being a CORDATE and that of being a RENATE.

Third, concepts are invoked to explain certain capacities and abilities of categorisation, classification and recognition. Thus, having the concept GREEN explains the ability to correctly recognise green things as green, and to differentiate green things from red and blue things. It also seems that this ability to correctly categorise things which fall under the extension of a certain concept is required in order to be said to possess a concept. If I, despite normal lighting conditions and having perfectly functioning colour vision, cannot correctly tell the difference between green objects and blue objects, then arguably I do not possess either of the concepts GREEN or BLUE.

Although there is certainly some degree of cohesion between these three roles, coming up with an adequate theory of concepts which accounts for all their functions has proven difficult. Due to the multiple functions which concepts are presumed to fulfill, there is a great number of desiderata which, allegedly, a fully satisfactory theory of concepts needs to meet. As an illustration, Jesse Prinz (2002) lists no less than seven features which a theory of concepts should account for; scope, intentional content, cognitive content, acquisition, categorisation, compositionality and publicity. Rather unsurprisingly, the result is that there is a plethora of theories of concepts in the current philosophical and psychological literature (see Prinz (2002) for an overview). Recently there has also been increased scepticism towards concepts, and many writers consciously avoid talk of “concepts” altogether. For example, Edouard Machery (2009) has advocated for concept eliminativism at least in *scientific* theorising, as concepts are not a natural kind and talk of “concepts” conflates multiple phenomena which we would do best to keep apart. Consequently, it isn’t clear whether we should aim to explain all the phenomena gathered under the heading of ‘concepts’ with reference to a singular theoretical entity.

In summary, contemporary discussions of concepts are highly complex and it is not clear that there is a unified subject matter to be addressed rather than several loosely related phenomena. Luckily, we need not settle on a correct and complete theory of concepts in order to

explain how “concepts” can be inconsistent. Instead, I will simply describe the relevant phenomenon I have in mind when I talk of “inconsistent concepts” and explain how this can underwrite the inconsistency approach. (Whether it is a mischaracterisation to speak of this phenomenon as inconsistent “concepts” is a metalinguistic dispute which I will not engage in here.)

2.2 What Are Inconsistent Concepts?

Although it appears that idea has been around for some time, the theory of inconsistent concepts is undoubtedly still in its infancy. Despite having received some attention in the most recent philosophical literature (see Scharp 2007; 2013; 2018; Greenough forthcoming; 2017; Cappelen 2018), the only positive account of inconsistent concepts is due to Kevin Scharp (2007; 2013). (Accordingly, I will base my discussion in this section mainly on Scharp’s account of inconsistent concepts). That being said, “incoherent concepts” also receive a mention in Sorensen (1992), Azzouni (2007) and Patterson (2009) discuss inconsistent semantic theories, and Eklund (2002) explores seemingly the same phenomenon but without mentioning ‘concepts’; instead he discusses inconsistent and incoherent semantic intuitions, competence, expressions, and Fregean senses.

I take this to be sufficient evidence that there is a relatively unified phenomenon here which warrants philosophical examination. Broadly construed, the phenomenon I am concerned with is the occurrence of inconsistencies and conflicts in the meanings of words in natural language, and how these inconsistencies may manifest in the associated semantic competences and intuitions. I propose, then, to use “inconsistent concepts” as a shorthand for referring to this phenomenon.

This use of the word ‘concept’ has some affinity with the second role of concepts mentioned above, on which concepts are the meanings, intensions, or Fregean senses of words. The reason why discussion of concepts is preferable to simply discussing words is naturally because words in themselves cannot be inconsistent; they are merely the vehicles or carriers of meaning. But the meanings or concepts which words express can be inconsistent (barring the radical externalist view that meaning is exhausted by reference), as will be explained very shortly. Thus, to say that TRUTH is an inconsistent concept is to say that the concept or meaning expressed by ‘truth’ is inconsistent. (Whether concepts in this sense are also constituents of

thoughts or what underlies certain capacities of categorisation is a further issue which need not currently be addressed.) We should also take care not to confuse concepts with the terms they are associated with. For example, some words, such as ‘bank’, may express more than one concept -- in this case the concepts RIVERBANK and FINANCIAL BANK. Similarly, there are sometimes multiple terms that express the same concept; arguably the English word ‘green’, and the French word ‘vert’, both express the concept GREEN.

Following Scharp (2013), I will take it that concepts have constitutive principles. A concept’s constitutive principles are the rules of application of the concept that are individuating for the concept in question. In other words, if these constitutive principles were to change, we would strictly speaking be dealing with a different concept. For example, (T-In) and (T-Out) are (presumably) constitutive of the concept of truth. If one of these rules of application were to change, we would no longer be dealing with our ordinary concept of truth, but a different truth-like concept.

As such, the application of concepts is a rule governed activity. To judge that a concept applies to some object or entity is to judge that the object or entity falls under the category or class of things which the concept represents. In order for a subject to qualify as possessing and understanding a concept, it needs to be possible to describe the subject as tacitly accepting a set of rules for its application, even though the subject might not consciously be aware of what the precise rules are (I will discuss concept possession in much more detail in section 2.3). For each concept, it should in principle be possible to describe its correct usage within a linguistic community with a set of rules that tell us when the concept applies and when it disapplies.⁹ These rules will also for the most part reflect the competent speakers’ semantic intuitions about how to use the concept (individual idiosyncrasies and errors may of course occur). The most central rules of application -- i.e. the rules of application which cannot change without changing the concept -- are its constitutive principles. The rules for applying a concept to some object are the same as the rules for applying the word expressing that concept to the object; i.e. if ‘blue’ applies to x then x falls under the concept BLUE. Likewise, if the concept BLUE applies to x, it is correct to call that object ‘blue’. (Hence, I will in the course of this discussion offer example rules of application for concepts that look like definitions of words, but strictly speaking, when we define words we are in fact defining the word’s *meaning* -- i.e. the concept.)

⁹ Following Scharp (2013), I use ‘disapplies’ as an antonym for ‘applies’. To judge that a concept *disapplies* to some object is to make the judgment that the object does not fall under the category which the concept represents.

For a concept to be inconsistent is for the concept to have constitutive principles (or minimally a set of rules of application) that produce inconsistency by disposing those who possess the concept to make conflicting judgements about when the concept applies and disapplies.¹⁰ How such inconsistent rules of application might look is perhaps best illustrated by an example. This one is borrowed from Scharp (2013, p. 36):

‘Rable’:

(1a) ‘Rable’ applies to x if x is a table.

(1b) ‘Rable’ disapplies to x if x is a red thing.

To clarify, the claim is not that the rules of application are inconsistent with each other, as there is no contradiction between (1a) and (1b) on their own. Rather, the inconsistency arises from the rules of application together with facts about the world. RABLE is an inconsistent concept because given that it is possible that there are objects which are both tables and red, i.e. red tables, its rules of application would imply that such objects both are and aren’t rables. In other words, the concept seems to have the implication that *there are no red tables*, since from the assumption that there are red tables we can derive a contradiction and thereby produce a reductio of “there are red tables”.

For illustration, assume that there is an x such that x is a red table. Together with (1a) we can then derive from our assumption that x is a rable. Similarly, from our assumption and (1b) we can derive that x is not a rable. So, x is a rable and x is not a rable. But that’s a contradiction. Therefore, assuming classical logic, there is no x such that x is a red table. Obviously our conclusion is false, seeing as there are red tables. An alternative way of formulating the the idea behind inconsistent concepts is thus to say that inconsistent concepts imply falsehoods (Scharp 2013), such as *there are no red tables*.

We can contrast the inconsistent concept RABLE with the consistent concept NON-RED-TABLE:

(2) ‘Non-red-table’ applies to x iff x is a table and x is not red

Important to note is that NON-RED-TABLE is not *just* RABLE made into a consistent concept, seeing as ‘rable’ both applies and disapplies to red tables, whereas ‘non-red-table’ clearly does

¹⁰ Some clarification is perhaps in order of what is meant by ‘inconsistency’ here. Standardly, inconsistency is defined as a property of a set of sentences -- but clearly, concepts are not sets of sentences. What is intended by saying that a concept is inconsistent is rather that the concept disposes those who possess it to make conflicting (or inconsistent) judgments.

not. An equal contender for a consistent replacement for RABLE would simply be TABLE, which applies to all and only tables, regardless of colour. RABLE must therefore be considered as a concept in its own right, importantly distinct from both NON-RED-TABLE and TABLE.

Additionally, we can note from the example that inconsistent concepts have rules of application and disapplication that overlap, so as to produce a set of objects to which the inconsistent concept simultaneously both applies and disapplies. For consistent concepts, the application set and the disapplication set are non-overlapping (Scharp 2013). For instance, NON-RED-TABLE applies to tables which are not red, and disapplies to everything else.

In addition, it is useful to distinguish inconsistent concepts from other kinds of defective concepts, such as partial, unsatisfiable and empty concepts. A partial concept is a concept which is only partially defined (Soames 1999, Scharp 2013). To give an example, consider the application rules of the following made up concept expressed by ‘freezy’:

- (3a) ‘Freezy’ applies to x if x is cooler than 5°C
- (3b) ‘Freezy’ disapplies to x if x is warmer than 10°C

For anything that is between 5°C and 10°C, it isn’t clear whether it is either correct or incorrect to say that it is “freezy” or “not freezy”, hence why FREEZY is partial. Thus, for a partial concept, instead of there being an overlap between the set of things to which the concept applies and does not apply, there is a set of objects for which it is undefined or underdetermined by the rules of application whether the concept applies or disapplies. Although partial concepts have a lot in common with inconsistent concepts, FREEZY is still a consistent concept because the set of objects to which it applies and set of objects to which it does not apply are disjoint.

Unsatisfiable concepts are also importantly different from inconsistent concepts. An example of an unsatisfiable concept would be the following (Scharp 2013):

- (4) ‘Squircle’ applies to x iff x is a square and x is a circle

Since nothing can be both a square and a circle at the same time, there is nothing which squircle could apply to; it is thereby unsatisfiable. There is obviously also a type of inconsistency present in the concept SQUIRCLE, since nothing can be both square and circular at the same time -- those properties are mutually exclusive. However, the concept itself is not inconsistent in the way in which RABLE is inconsistent, since it is perfectly clear when SQUIRCLE applies and when it disapplies -- it simply disapplies to everything. SQUIRCLE attempts to pull an object in two inconsistent directions at the same time, demanding that it be both square and circle. It is

obvious that no object could satisfy this demand. RABLE, on the other hand, merely pulls *its user* in two directions at the same time, demanding that its user both judges that the red table is a “rable” and at the same time that it is not a “rable.” While this is still a contradictory demand, it is perhaps less obviously so, seeing as there clearly are some objects to which the user can apply or disapply the concept without conflict. The crucial difference between an unsatisfiable concept and an inconsistent concept is thus that while unsatisfiable concepts place incompatible demands on objects in the world, inconsistent concepts place incompatible demands on its users (Scharp 2013).

It is also worth contrasting unsatisfiable and inconsistent concepts with empty concepts. Empty concepts are perfectly consistent concepts which place no incompatible demands on either the world or its users, but which simply happen to not apply to anything. For illustration, consider the concept expressed by ‘voon’:

- (4) ‘Voon’ applies to x iff x is a moon of the planet Venus

Because it so happens that Venus does not have any moons, VOON is an empty concept. Yet, there is nothing inconsistent or incoherent about supposing that Venus could have had moons, in which case VOON would not have been empty. Hence, VOON is both perfectly satisfiable and consistent, but, as a matter of contingency, does not apply to anything.

There are of course also many other ways in which concepts could be defective that I have not mentioned here. But hopefully I have said enough to make clear which specific conceptual defect is intended by saying that a concept is inconsistent, and how it differs from consistency, unsatisfiability, partiality and emptiness.

2.3 *Possessing Inconsistent Concepts*

2.3.1 *The Problem of Concept Possession*

The first objection to the claim that concepts can be inconsistent is that there is a problem when it comes to explaining how inconsistent concepts could be possessed; i.e. how one could have an understanding of and be competent with their use without being irrational. A common idea in the literature about what it is to possess a concept is that to possess a concept is to have a specific set of (implicit) beliefs -- perhaps one has to (implicitly) believe or accept the rules of application or some constitutive principles of the concept. However, it might be thought that this presents a problem for inconsistent concepts, because it is not clear how one could

characterise a concept as inconsistent without in the process ascribing inconsistent beliefs to those who possess the concept. In order to avoid this problem, we need to say a little bit about what it is to possess a concept, and why possessing an inconsistent concept does not result in ascribing inconsistent beliefs to those who possess inconsistent concepts.

To see why this is a pressing issue, consider what is perhaps the most prominent theory of what it is to possess a concept, namely, *conceptual role semantics*. According to conceptual role semantics, to possess a concept is to accept or (implicitly) believe its meaning constituting inferences. The meaning constituting inferences are related to the constitutive principles or rules of application mentioned above, as the meaning constituting inferences are generally taken to be individuating for the concept, and accepting the meaning constituting inferences is necessary for competence with the concept. If one does not accept these inferences, one arguably does not understand the meaning of the concept, but uses it with a different, idiosyncratic meaning (provided one succeeds in using it at all). The meaning constituting inferences for RABLE would plausibly be the following (these are directly adapted from the constitutive principles):

- (1c) If x is a table, then x is a rable.
- (1d) If x is red, then x is not a rable.

Hence, in order to possess the concept RABLE, conceptual role semantics says that one needs to accept inferences (1c) and (1d) above. But this creates a problem; if one also believes that there are red tables, it seems that one ends up with inconsistent beliefs. Consequently, it seems that anyone who possesses the concept RABLE must either deny that there are red tables, or be committed to an inconsistency. If we take it as a given that one ought not to believe contradictions, according to conceptual role semantics, anyone who possesses an inconsistent concept is as a result irrational.

However, it seems perfectly possible for subjects to possess and be competent with the concept RABLE, and not be irrational. For example, if I am competent with the concept RABLE while simultaneously acknowledging that RABLE is inconsistent, it seems that I would suspend judgment on either or both of (1c) and (1d) -- provided that I also admit that there are red tables. But even if I were to suspend judgment on (1c) and (1d), it seems that I would still possess the concept, so the theory of concept possession, as stated, cannot be correct. Modifying the account to talk of a *dispositions to believe* instead of belief simpliciter does not help much either, since if I possess the concept RABLE while believing that it is inconsistent I am also no longer

disposed to believe either (1c) or (1d), but it seems that I still understand, and am competent with, the concept. Therefore, in order to make room for inconsistent concepts, we need a different theory of concept possession than the one proposed by conceptual role semantics.

2.3.2 Scharp's Theory of Concept Possession

In order to solve this problem, Scharp (2013) argues that we should understand concept possession in terms of *entitlements* or what he calls *quasi-entitlements*.¹¹ On such a view, to possess a concept is to be entitled to the concept's meaning constituting inferences.¹² The notion of entitlement originates from Burge (1993), and characterises a special type of epistemic warrant to believe, which may be present also in the absence of evidence and is therefore weaker than justification. Entitlements, however, are defeasible, and only hold as long as the subject has no reason to doubt that which she is entitled to believe. To say that a subject is entitled to the meaning constituting inferences of a concept is thus to say that the subject has an epistemic warrant to accept the inferences, unless the subject has some reason to doubt them. Scharp's suggestion would as such be that to possess the concept RABLE is to be (quasi-)entitled to (1c) and (1d). According to Scharp, to be quasi-entitled to a claim is to be in such a condition such that one would have been entitled to the claim had one not had countervailing evidence (2013, p. 47). Accordingly, the idea seems to be that before I realise that RABLE is inconsistent, I am entitled to believe (1c) and (1d). Once I come to realise that RABLE is inconsistent, I lose my entitlement to believe (1c) and (1d), but I retain a quasi-entitlement to (1c) and (1d), in virtue of which I still possess the concept.

Scharp's suggestion is problematic for at least two reasons. Firstly, it seems that the notion of a quasi-entitlement doesn't actually do any explanatory work here. Unless a further analysis is offered of why the subject still retains a quasi-entitlement, it seems that 'quasi-entitlement' merely names the phenomenon to be explained, namely, how it is that one can remain competent with a concept which's meaning constituting inferences one rejects. With the notion of a quasi-entitlement in hand, we still have no better understanding of what concept possession consists in. Scharp seems to recognise this and makes a relatively brief remark that

¹¹ Due to considerations of brevity, I have decided to selectively discuss Scharp's positive account and my own proposal. But see Scharp (2013, 2.3) for further discussion and criticisms of other suggestions that have already been made in the literature.

¹² Scharp talks about 'constitutive principles' instead of meaning constituting inferences. But as far as I can tell, if one thinks of concepts as meanings, there is no significant difference between constitutive principles and meaning constituting inferences.

acquiring the (quasi-)entitlement requires a cognitive achievement, which takes the shape of acquiring an ability of some kind, presumably to use the concept in question in one's mental life. But if possessing a concept reduces to possessing an ability, we need further clarification on whether this ability in turn requires some knowledge-that or knowledge-how, and it is unclear what explanatory work entitlements are then meant to do in addition to the ability. Rather, it seems that Scharp has turned the order of explanation on its head; it should be concept possession which grounds the (quasi-)entitlements, not the other way around.

Secondly, it could potentially be problematic that invoking entitlements by means of concept possession incurs something reminiscent of *epistemic analyticity*. If Scharp is correct, it seems that in virtue of possessing concepts, we are suddenly now entitled to believe all kinds of things. In virtue of possessing the concept RABLE (provided I have not yet realised it is inconsistent), it seems that I am entitled to believe that "All tables are rables" (from (1c)). Moreover, if I am also entitled to believe that "No red object is a rable," it seems that I could infer that "There are no red tables." Hence, I would be a priori entitled to believe that "There are no red tables", just in virtue of possessing the concept RABLE. So, it follows from Scharp's account that merely by being competent with the concept RABLE I am entitled to believe substantive claims about the world. Presumably there is no upper limit to the number of inconsistent concepts that I could possess, and resultantly, on Scharp's account, I could in principle be entitled to a wealth of false a priori beliefs about the world.

If one accepts that there is such a thing as epistemic analyticity, there is of course no reason to think that this is problematic. However, epistemic analyticity is a fairly controversial matter, and Scharp himself acknowledges (2013, p. 54), and I agree, that Williamson (2007) has persuasively argued that nothing follows epistemically from mere conceptual competence. Williamson's arguments cannot be rehashed here, but a main takeaway is that it seems implausible to claim that a subject could be justified in believing a claim in virtue of merely understanding it (thus, in virtue of "mere" linguistic or conceptual competence). An entitlement is of course not yet a justification, but if one accepts that justification cannot flow from conceptual competence alone, one presumably needs to explain why entitlements are different and less problematic. Presumably entitlements would be less problematic because they are weaker, and thereby cannot ground knowledge. But if an entitlement does not amount to knowledge or a justified belief, it is rather unclear in this case in what way an entitled belief is superior to a completely unfounded belief. Usually when entitlements are invoked there is some

externalist story present in the background which explains the entitlement -- Scharp is yet to provide such a story. On the contrary, if one suspects that a large number of our concepts are inconsistent or otherwise defective, one has reason to believe that such an externalist story is not available. Hence, unless one has a plausible explanation of the entitlement, or is prepared to commit oneself to something resembling epistemic analyticity, entitlements are not plausible candidates for an account of concept possession.

2.3.3 An Alternative Proposal

In order to avoid the problems incurred by Scharp's theory of concept possession, I would propose that we think of concept possession as follows: to possess a concept is to master its rules of application. And to master its rules of application is to have implicit or explicit knowledge of the concept's rules of application. Most of the time, I suspect the knowledge of the rules will be implicit (or even some form of "knowledge-how"), and it won't necessarily be possible for the subject to articulate the rules on demand -- the subject just intuitively knows how to apply the concept. This explains how one might possess an inconsistent concept without realising that it is inconsistent, as one can master its rules without being consciously aware of them. Nevertheless, it should be possible to describe the linguistic behaviour of somebody who competently uses a concept as following a set of rules, even though the subject isn't necessarily aware of following them. Hence, to possess the concept RABLE is something like (implicitly) knowing the rules of application (1a) and (1b) mentioned above, or perhaps believing a modified version of (1c) and (1d), such that one merely infers that the concept or word applies to the object, rather than inferring something about the object, like below:

(1e) If x is a table, then 'rable' applies to x.

(1f) If x is red, then 'rable' disapplies to x.

The core idea is thus that to possess a concept is to have certain beliefs about the concept as a concept -- not about the thing the concept refers to or applies to.¹³

¹³Alternatively, we could denote the inferences required for concept possession as follows:

(1e*) If x is a table, then x is a "rable".

(1f*) If x is red, then x is not a "rable".

The distinction between applying a concept and inferring something about the object is thus between inferring that something is a "rable" and inferring that it is a rable.

Now, it might of course be argued that introducing this distinction cannot really help us here, because we are still attributing inconsistent beliefs to the subjects. If we let x be a red table, why should it be any less problematic for a subject to have the belief that ‘rable’ both applies and disapplies to x than it is for the subject to believe that x is a rable and that x is not a rable? While it certainly requires some explanation, I think it is actually a lot less problematic, for the reasons which will be outlined below.

Firstly, language is a human construct. Natural language has grown and evolved over a very long time and is in constant transformation. As the rules of natural language continually change, it is perhaps not so surprising to think that some inconsistencies among the rules have emerged along the way. And if a subject is competent with all these rules, there will inevitably be cases where the rules make incompatible demands on her.

Compare this to the EU Law. EU Law is another human construct which has evolved over several decades and now encompasses hundreds of articles. It is to be expected that in this extremely complex collection of rules there will be many cases for which the law might give incompatible recommendations; for example, that a certain act might be punishable by law under the right to religious freedom and at the same time that the very same act is protected under e.g. the freedom of the press. Laws are often conflicting or unclear; that is why lawyers have certain guidelines for how to interpret the law in cases where it conflicts. However, it wouldn’t be problematic to think that a lawyer who is highly knowledgeable about EU Law would when confronted with an especially tricky case initially come to the conclusion that *according to the law as it is written*, a certain act qualifies as both “illegal” and “legally protected.” That doesn’t mean that the lawyer has inconsistent or irrational beliefs -- it just means that she has knowledge of a conflicting system of rules.

Analogously, I don’t think it is any more problematic to think that a subject could simultaneously believe that it is correct to apply a concept and also believe that it is correct to say that the concept does not apply. What the subject believes is not that the object in question simultaneously both has and does not have a specific property -- that would of course be irrational. In other words, to judge that something falls into a certain category does not have to mean that one judges that it therefore has a specific property, other than the property of being in that category.¹⁴ That the subject has the semantic intuitions that a specific concept both applies

¹⁴ This distinction thus makes use of a sparse conception of properties, as opposed to an abundant conception (see e.g. Lewis 1986).

and disapplies to some object need not be more problematic or irrational than believing that according to EU Law, a specific article both applies and disapplies to a specific case.

An objection one might have at this point is that one cannot disassociate the judgment that a concept applies from the judgment that an object has a certain property. Thus, I cannot judge that ‘rable’ applies to x without judging that x is a rable. If this is right, we are in trouble -- the problem we thought we had solved simply comes back. To illustrate this point, assume that the following very plausible-sounding principle is true:

(Principle A₁) ‘F’ applies to x iff x is F

(Principle A₂) ‘F’ disapplies to x iff x is not F

If this principle is correct, it means that ‘rable’ applies to x iff x is a rable, and that ‘rable’ disapplies to x iff x is not a rable. So, from the belief that ‘rable’ both applies and disapplies to x, we will infer that x both is and is not a rable. And once again we have problematically inconsistent beliefs.

However, I think this is an argument in favour of rejecting principle A as false, rather than concluding that one cannot have inconsistent concepts without inconsistent beliefs. In fact, I think that rejecting principle A is precisely what anyone who comes to realise that one of their concepts is inconsistent does more or less implicitly. When we realise that RABLE is inconsistent, we don’t conclude that there is a special property of rablehood which red tables both have and do not have. Instead we conclude that ‘rable’ fails to reliably pick out a unique, coherent property in the world. We might conclude that there is no such thing as rablehood at all, although we have a concept for it. Or that it can only be captured by saying that something has the property of being a rable if it is such that people competent with the concept RABLE would point to it and say ‘rable’ (but if *that’s* the property, then there’s nothing problematic about thinking that something both is and is not a rable, as it is certainly conceivable that some of the population would so point, and the rest would not). The upshot is just that all “properties” or categories which it appears that we can construct or conceptualise in language do not succeed in tracking actual coherent properties which could be had by an object. Principle A should be rejected.

A related objection which deserves some attention is that some might take the preceding discussion to show that there cannot be inconsistent concepts because they cannot have meaning. Inconsistent concepts fail to consistently pick out anything in the world and so fail to be concepts at all. However, I think this objection is slightly too optimistic about our conceptual

capacities. And clearly ‘rable’ expresses a usable concept. For example, we can use ‘rable’ unproblematically to talk about non-red tables. ‘This looks like an expensive rable,’ I could remark about your white dining table, completely unproblematically and without conflict. ‘That’s the ugliest rable I’ve ever seen,’ I could say about a red table spotted at a flea market, and probably you would accept my utterance, because it isn’t completely obvious that I’ve made any linguistic violations and probably you can still correctly interpret what I mean. So clearly ‘rable’ can have meaning. And if ‘rable’ has meaning, then arguably RABLE is a concept.

In summary, then, possessing inconsistent concepts is not necessarily problematic if we think of concept possession as having knowledge of rules. In other words, being competent with a concept is to have knowledge about the concept itself as a concept -- not about the things it refers to or picks out in the world (if any).

2.4 Error Theory and Externalism

The idea that there are inconsistent concepts has recently been strongly criticised by Herman Cappelen (2018), who relies an externalist metasemantic framework to defend an error theory about inconsistent concepts and argues that the occurrence of inconsistent concepts amounts to an illusion.

The main motivation behind this view appears to be that Cappelen doesn’t want to commit himself to the existence of “concepts.” As should be evident from my earlier remarks, I have a lot of sympathy for reluctance to commit oneself to “concepts”, but I think the way that I have spelled things out so far should make clear that we do not strictly need to talk of “concepts” to address the phenomena at hand -- it’s just a very convenient shorthand for referring to “the aspect of meaning of words not exhausted by their reference.” Accordingly, it seems to me that many of the points raised by Cappelen are entirely compatible with there being inconsistent concepts, and in fact many of the externalist concerns cited by Cappelen in favour of his error theory help to motivate and explain the phenomenon of inconsistent concepts -- at least in the way that I have outlined them here. I shall spell out in more detail why that is in what follows. I will group Cappelen’s criticisms together into two main lines of objection.

Cappelen’s first line of objection is that when we think that a concept F is inconsistent, it is in fact our beliefs about fs that are inconsistent, and not the concept. He writes:

According to the kind of externalism that I endorse, there’s no cluster of beliefs (or dispositions to endorse) that are analytically true or are required for being a competent user of

a predicate. Speakers can be deeply committed to *all kinds* of false beliefs about Fs and still be competent users of 'F' and use 'F' to say false things about Fs.

(Cappelen 2018, p. 86)

The first point is thus that according to Cappelen, no specific beliefs are required to be competent with a predicate 'F'. The second point is that there is no limit to the false things which subjects may believe about Fs, and they can still use 'F' to talk about Fs. For example, on Cappelen's view I could believe that all pencils are aliens and still use 'pencil' to talk about pencils. The third and final point which Cappelen makes on this topic is that it is inconsistent *beliefs* about Fs that explain the appearance of an inconsistent concept, while the concept expressed by 'F' is in fact perfectly consistent.

Cappelen's first point is clearly in direct conflict with my favoured theory of concept possession. Unfortunately, Cappelen doesn't give us an alternative theory of concept possession (presumably because he denies that there are concepts). But what can be noted is that it is rather mysterious what exactly concept (or predicate) possession could consist in if there are no beliefs or dispositions required to possess a concept or predicate. Surely *something* must be required which explains why it is that I don't understand or am competent with all the world's concepts. Perhaps this need not be a belief, although I find that theory most plausible. Moreover, endorsing a different theory of concept or predicate possession need not be incompatible with what I have said so far about inconsistent concepts.

In regards to the second point, there is no real conflict between what Cappelen says and my theory of concept possession. On the theory I favour, there are no restrictions on what one may believe about Fs and still possess the concept expressed by 'F'. All that is required for me to possess the concept CAT is that I have certain beliefs about the concept CAT -- i.e. beliefs about when it applies and when it disapplies. Without knowledge of the rules of application of 'cat', I couldn't fully understand others when they speak of 'cats'. But that does not prevent me from believing that cats are robots from the planet Mars, since this is a belief about the entities which 'cat' refers to, about cats -- not about CAT as a concept, or about when 'cat' applies or disapplies.

Interesting to note is that Cappelen admits that individual subjects could have inconsistent *conceptions* of a concept, although the concept itself is not inconsistent. That is because he thinks that it is inconsistent conceptions of concepts that explain the appearance of inconsistent concepts. But if one admits that there could be inconsistent conceptions of

concepts, then surely there could also be inconsistent concepts -- that would just amount to these inconsistent conceptions being shared by a community of individuals.

Lastly, why are inconsistent concepts not explained away by inconsistent beliefs about the thing that the concept refers to? This is a trickier question, but an important one. I think we need the notion of inconsistent concepts to explain certain conflicting semantic intuitions which may be present although we have no object-level inconsistent beliefs about the thing which the concept refers to. Take for example the concept of knowledge, which might be thought to be inconsistent. Think of the standard sceptical puzzle where the subject doesn't know that she is not a brain in a vat but still knows that she has hands. Now in this case we are inclined to say that she knows that she has hands, but also that she doesn't know that she has hands, because she cannot rule out that she's a brain in a vat (in which case she wouldn't have hands). Our intuitions conflict about whether it would be correct to say that the subject "knows" or does not "know". But clearly we have no inconsistent beliefs about the subject's mental states. We have perfectly consistent beliefs about the etiology and justification of her beliefs about her hands and the etiology and justification of her belief about whether she's a brain in a vat. We have completely consistent beliefs about the probability or credence she attaches to the proposition that she has hands or the proposition that she's a brain in vat. But still we feel the semantic pull of our concept of knowledge to make incompatible judgments regarding whether she "knows". For what it's worth, we might even have particle for particle perfect knowledge of her brain state and how that is related to her hands and the possibility of being a brain in a vat but still find ourselves unable to say whether she "knows"! This phenomenon, I think, is much better diagnosed by appealing to inconsistent concepts than inconsistent beliefs.

The second line of objection pursued by Cappelen stems from externalist considerations of "metasemantic messiness". For example, Cappelen notes that "a term can be introduced through reference-fixing descriptions that impose incoherent or inconsistent conditions on the referent" (2018, p. 87). He also recognises that there can be multiple or conflicting baptisms, and that baptisms in general are a very messy ordeal. Further, reference changes may occur which has as a result that reference fails, or which leaves us inclined to think that there are multiple equally good candidate extensions for the term. Sometimes metasemantic conflicts and tensions occur which can be resolved in more than one way, both which seem equally viable. Sometimes we also defer to experts in fixing references, which can make reference messy and non-transparent. All

these considerations taken together, Cappelen thinks, add up to a state of metasemantic chaos which can explain away the illusion of inconsistent concepts.

In contrast to Cappelen, I think all these considerations of metasemantic messiness in fact help us explain why there are inconsistent concepts, and how they come about. What I have in mind when speaking of inconsistent concepts is precisely inconsistent rules of application for concepts that result from these kinds of messy metasemantic phenomena. Take for example the metasemantic tensions which can be resolved in more than one way, which were mentioned by Cappelen. This is precisely what we would expect if there were inconsistent concepts. Let's go back to our by now familiar concept RABLE. When subjects who possess the concept RABLE encounter a red table, they will feel a metasemantic tension. Either they can decide to call it a "rable," and privilege rule (1a), or they can decide to refrain from calling it a "rable" by privileging rule (1b). Both are equally good options. There are thus two ways of resolving the metasemantic tension, either by collapsing the meaning of 'rable' into the meaning of 'table', or by collapsing the meaning of 'rable' into the meaning of 'non-red-table', and there are no clear guidelines here for how the tension should be resolved. Different individuals might resolve it differently, creating their own consistent conceptions of the concept RABLE. But suppose that there is no clear majority, and both uses of 'rable' persist. Then both rules (1a) and (1b) will remain part of the concept RABLE, since individuals will need to have knowledge of both of the rules in order to communicate effectively with each other. One cannot just suppose that one's interlocutor has resolved the tension in the same direction as oneself and start speaking of "rables" presupposing that 'rable' definitely applies even if the table is red or vice versa. That's a recipe for misunderstandings. Moreover, if meaning-facts supervene on community-wide use-facts, then it seems that given that people have resolved the tension in different ways, the inconsistency will remain in the meaning of 'rable', since the use-facts are still conflicting. Hence, it is possible for a concept to be inconsistent although most individuals might have developed their own consistent conceptions of the inconsistent concept and thereby resolved the tension. These considerations, I believe, speak more in favour of there being inconsistent concepts than of Cappelen's error theory.

Finally, we might ask, is the phenomenon of inconsistent concepts perhaps still incompatible with externalist metasemantics more broadly? I will concede that, because the view requires that there are rules of application or meaning constituting inferences or principles, it does indeed seem incompatible with the most radical forms of externalism, on which there are

no such rules or application conditions. I cannot refute such a view here -- an extended discussion of metasemantic externalism certainly falls outside of the scope of this thesis -- so I will for my purposes here simply assume that such a radical externalist view about meaning is false.

To sum up, it seems to me that Cappelen's error theory doesn't actually succeed in debunking the occurrence of inconsistent concepts; he merely describes the phenomena using different language to avoid committing himself to "concepts". The only substantial point of disagreement remaining here, then, appears to be whether it is useful to talk about "concepts" and whether concepts have rules of application at all. But I think I have said enough about this matter already.

3.

THE INCONSISTENCY APPROACH TO EXISTENCE

What is existence? Or rather, what does ‘existence’ *mean*? The inconsistency approach to ontology proposes that ‘existence’ expresses an inconsistent concept. In the previous chapter, I outlined a preliminary theory of what inconsistent concepts are and how they might be possessed. This chapter defends the inconsistency approach to existence, and shows how it can help us make sense of ontological debates. I will suggest that ‘existence’ is a little bit like ‘rable’; its rules of application lead us to in some cases judge that the concept both applies and disapplies. To flesh out the theory that the concept of existence is inconsistent, I will start by illustrating how the inconsistency approach can be motivated by one of the classic paradoxes of existence; the problem of non-being. I will then offer a crude example conceptual analysis of the rules of application of ‘exists’, and respond to some initial objections. The aim is not to defend the specific analysis per se, but the idea that the predicate contains an inconsistency. In section three I will develop the metaontological advantages of the inconsistency approach, and argue that the approach improves on Hirsch’s, Thomasson’s and Chalmers’ deflationary approaches discussed in the first chapter. The final section answers some residual objections to the inconsistency approach.

3.1 The Problem of Non-Being

A *prima facie* case for thinking that existence is an inconsistent concept comes from what is perhaps the most well-known philosophical puzzle of existence; namely, *the problem of non-being*. The problem dates all the way back to Plato’s *Sophist*, and is sometimes referred to as ‘Plato’s Beard’. Quine sums it up in the following catchphrase; “Non-being must in some sense be, otherwise what is it that there is not?” (Quine 1948, p. 21).

More elaborately, the problem of non-being is the problem posed by the ontological status of entities which can be meaningfully said to “not exist.” For instance, in some contexts it

seems perfectly correct and truthful to make negative existence claims, such as ‘Pegasus does not exist’, ‘Hogwarts does not exist’ or ‘There is no such thing as phlogiston’. However, for a sentence to be meaningfully asserted about e.g. Pegasus, it appears that there must be something which ‘Pegasus’ refers to, designates, or is about -- if not an actual divine winged horse, then perhaps we are talking about the idea of Pegasus, or a psychological entity of sorts. Yet, if there is then indeed *something* which ‘Pegasus’ refers to, then it seems that the original sentence ‘Pegasus does not exist’ must be false, because otherwise it is puzzling how the assertion could be meaningful at all. The same argument holds not just for names or singular terms, but also for general terms like ‘phlogiston’ or ‘dragons’. It seems that those terms, too, must refer, or be about something, in order for the terms to be meaningful. But if there is something which ‘phlogiston’ refers to, then it seems no longer true to say that “phlogiston does not exist” (Quine 1950; Thomasson 2009).

Consequently, for negative existence claims to come out as both meaningful *and* true -- as we frequently take them to be in ordinary speech -- they seem to require that there is something which does not exist. Some philosophers have taken this to indicate that being and existence can come apart, such that there can *be* things which do not *exist*. Meinong (1904/1981) famously argued that existence is best conceived of as a property which entities may have or may not have in addition to being, and that not everything there is exists. On the Meinongian picture, there are thus different modes of being; there are non-existent entities in addition to the existent entities, which are sometimes said to have *subsistence* instead of existence. There are also contemporary advocates of views in this vein, see for example Parsons (1975; 1980; 1982), or McDaniel (2009; 2010a; 2010b) who defends a view on which there are several modes of being.

The opposing view, frequently referred to as the Quinean picture due to his seminal exposition of the view in his classic paper ‘On What There Is’ (1948), takes existence to be univocal. On this picture, there is only one mode of existence or being which all expressions such as ‘exists’, ‘there is/are’ and ‘some’ express, and this can be adequately accounted for by the existential quantifier ‘ \exists ’. Quine argues that the answer to the question ‘What is there?’ has to be ‘Everything’, and that anything which does not exist cannot *be* at all. There is as such no distinction to be made between being and existence; Quine takes it to be a tautology that *everything there is exists*. Further, Quine thought that quantification entails ontological commitment; quantifying over something automatically entails that the entity exists. “Superfluous” ontological commitments to universals, fictional entities, and the like are to be avoided by paraphrasing

statements about these in natural language into a maximally reductive scientific language which might then be formalised, such that we are only committed to the existence of those entities which we must quantify over in our best (simplest) scientific theories. The Quinean view of existence is now commonly taken to be the standard view (see Crane 2012; for a contemporary articulation of the Quinean view see van Inwagen 1998; 2009).

The history of the debate surrounding the problem of non-being has many more sides to it than simply Meinongians versus Quineans, and a much richer history and literature than I could hope to do justice to here. The purpose of this short summary is to illustrate how the central dividing line in the debate -- whether there is or isn't a distinction between being and existence -- embodies what plausibly appears to be an inconsistency in the concept of existence. Conceiving of existence as an inconsistent concept immediately helps us make sense of this debate.

3.2 Rules of Application For 'Exists'

Let's return to the original paradox created by negative existence claims; namely, that non-existent entities must be *something* in order to be meaningfully talked about at all. A way of explaining how this paradox comes about is to say that 'exists' has inconsistent rules of application. Without committing ourselves to the claim that these are the exhaustive or definitive rules of application, we can note that both of the below are plausible candidates as rules of application for 'exists';

- (5a) If there is an x, then 'exists' applies to x
- (5b) If x is fictional, then 'exists' disapplies to x

Together with *Principle A* discussed in the previous chapter, it appears that (5a) and (5b) would license the following inferences:

- (5c) If there is an x, then x exists
- (5d) If x is fictional, then x does not exist

Now note that these rules seem to have the implication that *there are no fictional entities*. To see how this comes about, assume for reductio that there is an x such that x is a fictional entity. From our assumption and (5c), we can then infer that since there is an x, x exists. Similarly, from our

assumption and (5d), we can infer that since *x* is fictional, *x* does not exist. So *x* exists and *x* does not exist, which is a contradiction. By reductio, there are no fictional entities.¹⁵ In other words, if (5a) and (5b) are part of the concept of existence, then existence is an inconsistent concept, because its rules of application together with Principle A imply a falsehood.

The fact that there are fictional entities now poses a problem for anyone who possesses this concept of existence; in virtue of possessing the concept, and simultaneously observing that there are such things as fictional or other non-existent entities, one is led into a paradox. Granted that it is irrational to believe contradictions, the tension must be somehow resolved.

Without rejecting fictional entities, or the concept of existence altogether, there are two principal ways of resolving the tension. They both involve privileging one of the two conflicting rules of application in order to make the concept consistent. The first way to go would be to privilege rule (5a) in favour of (5b), and accept that fictional entities also do exist. Call this alternative A. The second alternative is to privilege (5b) over (5a), and reject the idea that everything there is exists, such that it can truly be the case that there are some things which do not exist, in addition to those that do. Call this alternative B.

Alternatives A and B clearly map onto the Meinong-Quine dispute above. In fact, it would appear that the Quineans are going for alternative A, while the Meinongians prefer alternative B. When Quine asserts that everything there is exists, he is asserting his unwavering endorsement of rule of application (5a), and accordingly he has to either reject (5b) or reject the claim that we must quantify over fictional entities. Meinong on the other hand, in accepting a distinction between existence and being, is privileging rule (5b) and rejecting (5a), as he believes that not everything there is also exists. The essence of the debate over the problem of non-being can thus be crudely accounted for by appealing to the theory that ‘exists’ expresses an inconsistent concept, and that the different solutions to the problem are different ways of resolving the inconsistency in the concept by privileging different rules of application, in order to make the concept consistent.

As I already stated, there is undoubtedly more to the concept of existence than these two rules. In order to come up with a fully adequate and complete set of rules of application for ‘exists’ which can fully characterise its use, we would certainly have to add some more rules and/or qualify the rules that are already there. However, I do not think it is all too implausible to

¹⁵ This line of reasoning of course presupposes that we haven’t yet realised that ‘exists’ is inconsistent, and that we have not yet rejected Principle A, as discussed in the previous chapter.

think that the two rules (5a) and (5b) above are both somehow baked into our undoubtedly “messy” concept of existence. As long as they are both somehow part of the concept, chances are the concept will turn out to be inconsistent. It could of course still be the case that once we have added the appropriate qualifiers or additional rules, it will turn out that EXISTENCE is not an inconsistent concept at all. Perhaps the fact that it seems inconsistent will turn out to simply be an artifact of our incomplete analysis -- that is for time to tell.

In order to uncover the exhaustive and definitive rules of application for the concept of existence, then, we need to do more, and more rigorous, conceptual analysis. On a more traditional picture of philosophical methodology, the aim of conceptual analysis is to find the *necessary and sufficient* conditions for a concept. However, if the concept is inconsistent, it will not be possible to find a unique set of conditions that are both necessary and sufficient. The best we can hope for is to find a set of rules of application which account for the complete use of the term and accurately reflect its meaning. If the aim is to find a descriptively accurate analysis of how the concept of existence actually works, it might also be beneficial to empirically investigate ordinary people’s semantic intuitions and use of the concept, in order to reliably and scientifically establish what its actual rules of application are. This could help us avoid unnecessary bias, considering that a small homogenous group of philosophers’ intuitions about the rules of application of a concept are unlikely to be representative of how the linguistic community at large uses the concept. In other words, to find an adequate descriptive analysis of existence it seems appropriate to turn to *naturalised conceptual analysis* (see Machery 2017).

Not to be overlooked is the huge number of conceptual analyses of existence already out there in the philosophical literature; certainly many more than any individual philosopher has the time to thoroughly consider. (Meinong’s analysis of existence as a property, and Thomasson’s *schema E* from the first chapter, would be examples of such analyses.) The crucial point, however, is that philosophers have almost exclusively looked for consistent conceptual analyses. Unfortunately, no consensus appears to have been reached so far on what the correct analysis of ‘exists’ is. I think this is evidence in favour of the inconsistency approach to existence. If ‘exists’ expresses an inconsistent concept, this would explain why nobody to date has been able to find an uncontroversial and fully satisfactory analysis of ‘exists’ that does justice to all of our pre-theoretical intuitions about existence and preserves all of its ordinary use.

This is because it is, as a matter of fact, impossible to give accurate and complete *consistent* analyses of inconsistent concepts. Because all consistent analyses involve rejecting one of the

conflicting rules of application of the concept, all consistent analyses will invariably have some counterintuitive implications or fail to account for some aspect of the concept. As a result, if the concept of existence is inconsistent, a consistent analysis will never be able to satisfy all of our pre-theoretic semantic intuitions about how the concept is to be correctly applied. Our failure as of yet to find an uncontroversial analysis of existence could therefore be seen as symptomatic of its inconsistency. The bottom line, then, is that there could still be a consistent and fully satisfactory analysis of the concept of existence out there waiting to be found -- but as things stand, it seems improbable.¹⁶

It also seems to be a live possibility that the inconsistency in the concept could be more accurately described by a very different set of rules of application, rather than rules (5a) and (5b) which I have suggested here. In fact, it is blindingly obvious that rules (5a) and (5b) are highly incomplete (although it also seems plausible to me that they both could somehow be part of the exhaustive and definitive set of rules). I therefore want to emphasise once again that I am by no means claiming that these are exhaustive or definitive rules, and nothing hinges on this -- they are primarily intended to illustrate the principle behind the idea and its explanatory potential. So long as there is an inconsistency present, metaontological explanatory benefits can be reaped.

3.3 Metaontological Advantages

The view that ‘exists’ is an inconsistent concept also has many metaontological advantages when it comes to explaining what has gone wrong in ontological debates, and why they still persist. In chapter one, I argued that none of the deflationary accounts succeed in adequately explaining what is going on in ontological debates. I believe that the theory that EXISTENCE is an inconsistent concept can better account for the nature of ontological debates, and why they have continued to puzzle us, than the other metaontological views. At the same time, some insights from each of the deflationary views can be preserved.

Firstly, it should be noted that the inconsistency approach retains a certain affinity with Chalmers’ anti-realist account. In broad strokes, Chalmers’ first-order semantic analysis of the absolute existential quantifier as a defective concept is partly in agreement with the inconsistency approach. The inconsistency approach develops the theory one step further by hypothesising that the defect in the concept of existence is an *inconsistency*. An important difference between

¹⁶ A related issue is of course that many philosophers, in doing conceptual analysis, tend to offer not purely descriptive analyses but also normative proposals for what a given concept *should* mean. More on this below.

Chalmers' account and the inconsistency approach is that, on Chalmers' view, it is only the philosophers' concept of absolute existential quantification that is defective while our ordinary concept of existence is non-defective. On the inconsistency approach, however, there is no such distinction -- there is one central concept of existence, and that concept is inconsistent. The reason why the inconsistency does not cause ordinary existence assertions to *seem* problematic is that these are still correct applications of the concept of existence, in the sense of obeying its rules of application. In addition, as long as both speaker and interlocutor possess the same inconsistent concept communication can still unproblematically take place. Hence, the problems don't become salient until we start reflecting on what existence is, or what 'existence' means, and realise that our semantic intuitions about when to apply and disapply the concept are inconsistent.

Moreover, the defect Chalmers has in mind seems to be rather something like the partial concepts mentioned in the previous chapter, whose rules of application leave a set of objects to which it is undefined whether the concept applies or disapplies. In other words, he seems to have in mind some sense in which the meaning or truth conditions of existence claims are *underdetermined*. The inconsistency approach, on the other hand, does not need to postulate any underdetermination, seeing as it might *both* be a correct application of the concept of existence to claim that "numbers exist" and to claim that "numbers don't exist". In my opinion, this better explains why ontological disputes seem irresolvable than underdetermination, as according to the application conditions for exists, it is *both correct* to say of certain objects that they do exist and that they do not exist. Each party can therefore keep insisting that *they* got it right; when the universalist says that "numbers exist", and the nihilist replies that "numbers don't exist", they are both using our concept of existence correctly. In a sense, then, both sides are partly correct. Contrastingly, on Chalmers view, it would seem more natural to say that *nobody* is correct.

There are thus important differences between mine and Chalmers' views, but still some similarities. As a consequence of the similarities, the inconsistency approach also enjoys what I argued was the main virtue of Chalmers' view -- that we can hold on to the idea that serious ontologists are attempting to make substantive, non-trivial claims about the world. On the inconsistency approach, just like on Chalmers' approach, we do not need to maintain that the ontologists are making trivial claims in their own languages, or merely arguing about how language should be used. The serious ontologists are all competent with our ordinary concept of existence and are trying to use this very same concept to make interesting claims about the

world. We can therefore avoid the problems of uncharitably interpreting the parties to ontological debates encountered by Hirsch and Thomasson.

Significantly, though, the inconsistency approach also allows us to incorporate some of the insights from the verbal disputes view and the metalinguistic negotiations account. As a matter of fact, on the inconsistency approach, *it is to be expected* that ontological debates sometimes end up being merely verbal or develop into a metalinguistic negotiations, due to the inconsistency in the concept of existence. The inconsistency approach can thus *explain* why it comes about that ontological disputes are linguistically defective in just the way that these other views argue.

According to the inconsistency approach, what in effect happens when ontologists begin to theorise, is that they come to detect an inconsistency in the concept of existence. In some cases, the philosophers have conflicting semantic intuitions about when the concept applies and disapplies; they detect a (meta-)semantic tension. As was explained in the previous chapter, in order for the subject to maintain a consistent set of beliefs, this tension must be resolved one way or the other, by effecting a subtle change in the rules of application of the inconsistent concept to produce a consistent concept. Hence, when the subject adopts a consistent set of rules of application for the concept, a shift in speaker's meaning occurs, and the subject goes from speaking ordinary English to speaking in their own language. As different individuals might resolve the semantic tension differently, it is entirely possible for two individuals to end up talking past each other as they, in effect, are using the same expression for different concepts. The inconsistency approach therefore has a natural explanation to how and why the shift in language (which Hirsch fails to explain) comes about. The philosophers aren't confusing themselves into speaking in their own language -- the shift in language is a result of resolving contradictions and becoming *less* confused!

In addition, the inconsistency approach can explain why pointing out that the dispute is "merely verbal" does not make the dispute go away. For example, if each side has settled on their own consistent conception of existence, which potentially succeeds in reliably tracking some property or structure in the world, it would indeed seem that each party could agree that the other speaks the truth in their own language. However, even though each side to the dispute then does mean something different by 'exists', and could in theory agree on this -- the parties are still disagreeing over which conception of existence best fits our commonsense concept of existence. Accordingly, each party could have the intuition that they got it right and the others

got it wrong, because they think their analysis better reflects *the* concept of existence. We can thus maintain that the disputants are not *just* talking past each other, but are in a sense still talking about the very same thing -- our concept of existence.

In effect, these participants are engaging in a metalinguistic dispute about the meaning of the concept of existence. Sometimes this might be a purely *descriptive* metalinguistic dispute (see Plunkett 2017), as the participants might be convinced that their conception of existence is a correct characterisation of our ordinary (inconsistent) concept of existence -- that they have uncovered what it meant all along. In other cases, it will be a *normative* metalinguistic dispute -- what we previously referred to as a *metalinguistic negotiation* -- where the participants are arguing over how we *ought* to conceive of existence. This is frequently a feature of revisionist ontology, where the revisionist ontologists may sometimes acknowledge that their conception of existence does not match the ordinary concept, but they still implicitly or explicitly maintain that their notion of existence is more interesting, or makes more sense, or for some other reason is what we really should be talking about. As highlighted in chapter one, though, these revisionist ontologists are generally not motivated by some normative or practical concerns, but usually insist that we use their conception of existence because it better manages to get at some feature of the reality which they are attempting to describe -- or plainly because the other notion leads us into contradiction. As such, even if some ontological disputes are in some sense “merely verbal”, the inconsistency approach explains why accepting that the dispute is verbal does not succeed in resolving the dispute.

Crucially, though, the inconsistency approach can also accommodate the strong intuition that the ontologists have that they are not confused about language, but are addressing substantial metaphysical problems. If each party to the debate has settled on their own consistent conception or analysis of the concept of existence, and are using this consistent conception to theorise about ontological matters, they might very well succeed in having consistent and substantial ontological thoughts which are not just about language. Hence, even if our ordinary concept of existence fails to reliably track some unique property or structure in the world, it is possible that some consistent conceptions of the concept of existence do. However, on the assumption that meaning supervenes on community wide use-facts, when ontologists make ontological claims in a public dispute, the concept of existence assumes its ordinary inconsistent meaning, which has the result that their claims might fail to be interpreted as the substantial claims they intended to convey (or at best, that only those who have already accepted their

particular analysis can interpret them correctly). As a result, even if ontological disputes turn out to frequently be linguistically defective, this is compatible with the idea that ontological debates have the *potential* to be about substantial matters pertaining to the nature of reality, once the linguistic obstacles and confusions have been cleared away. The approach thus suggests a natural way of correcting and eventually resolving defective ontological disputes; via conceptual engineering.

In summary, then, there are many metaontological explanatory advantages of thinking of the concept of existence as inconsistent. Firstly, it allows for a charitable interpretation of what ontologists are up to, as it vindicates the ontologists' self-conception as debating factual and non-trivial claims about the world, while also allowing the disputants to be genuinely disagreeing. Secondly, the inconsistent concepts approach does justice to -- and substantiates -- the deflationary intuition that there is confusion rooted in language at the heart of the debate, and explains why it is that ontological disputes sometimes end up being "merely verbal" or metalinguistic. Finally, the view explains why ontological debates are so persistent and seemingly irresolvable, as consistent conceptual analyses of existence are inevitably never going to fit the bill, despite being partly correct.

3.4 Objections to the Inconsistency Approach

Thus far I have argued that the inconsistency approach does a better job than its deflationary rivals in accounting for ontological disputes. There are, however, some residual objections to the inconsistency approach, which I will attempt to address in this section.

Firstly, what about other realist views? So far I have argued that neither Hirsch's, Thomasson's, nor Chalmers' deflationary metaontological accounts offer satisfactory explanations of what is going on in ontological disputes. Additionally, I argued in chapter one that none of the accounts succeed in deflating ontological debates, as they do not succeed in blocking the possibility of asking and answering substantial and non-trivial ontological questions. Hence, we should be realists about ontology. But aren't there some rival non-deflationary realist accounts which might be superior to the inconsistency approach?

As far as I'm aware, the inconsistency approach isn't directly in conflict with any realist account in the current literature. That is because those accounts do not generally attempt to explain the phenomena which the inconsistency approach aims to account for -- namely, why it

is that ontological disputes sometimes appear to be merely verbal or metalinguistic disputes, or why it might seem that the concept of existence has no determinate meaning. Realists sometimes acknowledge that *some* ontological disputes could be linguistically defective in this way, but that this does not threaten the possibility of doing serious and substantive ontology, which is a claim I'm fully in agreement with. Usually, however, the realists stop right there, which means that the inconsistency approach is mainly in competition with the deflationary views. The inconsistency approach would therefore only be in conflict with a realist view which positively maintains that the concept of existence is completely non-defective (or has some other defect that is not inconsistency).

To my knowledge, the only account which comes close to making any such a claim is Theodore Sider's ontological realism (2009; 2011). Sider argues that 'existence' carves at the joints of reality -- i.e. that the concept of existence tracks some natural property or structure -- which would seem to be in direct conflict with my claim that the concept of existence is inconsistent and fails to consistently or uniquely track anything at all.

To go into some detail about Sider's view, Sider believes that the world has a natural *structure*, and the aim of metaphysics is to discover what this structure is. "Structure" in Sider's technical sense is intended to refer to those features in virtue of which the world is not just "an amorphous blob" which we can carve up as we like by means of language. In other words, it refers to those features of the world which for example terms like 'green' and 'blue' seem better at picking out than 'grue' and 'bleen'. Those terms that perfectly succeed at reflecting reality's structure, or "carve nature at its joints", are called joint-carving terms. Joint-carving terms are epistemically better than non-joint-carving terms. In addition, natural properties and structures are so called *reference magnets* -- they attract the reference of our words. Crudely, the thesis of reference magnetism says that if there is a natural property or structure in the vicinity of how a term is generally used (i.e. its use maps on quite well, but not perfectly, to some natural property or structure), then that natural property or structure is generally the best candidate for the reference of that term. Reference magnetism could then for example explain why 'water' refers to H₂O (and always has), instead of just referring to any other clear, drinkable liquid.

According to Sider, 'existence' is a joint-carving term, which refers to reality's quantificational structure. Sider argues that there is a natural property of **existence** (I will follow Sider and use **bold** to indicate this (hypothetical) natural property), which, due to reference magnetism, our concept of existence tracks or refers to. This means that ordinary use, semantic

intuitions, and conceptual analysis are not reliable guides to what ‘existence’ actually means. Sider writes:

[I]f **existence** is what we mean by ‘there exists’ because of its reference magnetism, not because of its fit with ordinary use of ‘there exists’, then conceptual analysis needn’t be a guide to the truth values of English statements of existence [...]

(Sider 2009, p. 410)

And further:

[My] arguments against deflationism assume that the ‘force of reference magnetism’ is strong enough to outweigh a failure of **existence** to match the use of ‘there exists’ [...] Put less metaphorically, they make an assumption about the true theory of content determination: that this theory weights naturalness heavily enough to overcome any mismatch there may be between **existence** and the use of ‘there exists’. Whether this assumption is correct depends on the strength of the magnetic force (i.e., the relative weights of naturalness and use in the true theory of content), and also on the degree to which the sentences whose use **existence** does not match are meaning-constitutive.

(Sider 2009, p. 411).

This is thus a type of externalist objection, which I discussed in the previous chapter. I conceded there that the theory of inconsistent concepts is not compatible with radical externalism, and I cannot refute radical externalism about meaning here. But in the case of existence, I think problems loom large -- regardless of the correctness of radical externalism.

Firstly, I think there is a good chance the inconsistent rules of application of EXISTENCE are “meaning-constitutive” as Sider puts it, i.e. they are individuating for the concept of existence (if the rules were different, it would be a different concept). If that is the case, then there cannot be a unique natural property of existence which really matches our concept of existence, as no unique natural property could satisfy the “meaning-constituting” inconsistent rules of application. Further, considering that we currently don’t have a very good idea of what this natural property of **existence** would be, or whether there is only one such natural property in the vicinity rather than multiple good candidates, it seems that, even under the assumption that the concept is not inconsistent, it is very uncertain whether ‘existence’ really is a joint-carving term.

In addition to this descriptive argument, Sider makes a related normative point; even if it turns out that ‘existence’ in ordinary English does not refer to **existence**, **existence** is what we *should* mean by the term. If our ordinary use of ‘there exists’ fails to pick out the natural property or structure of **existence**, the ontologists can introduce their own language, *Ontologese*, and stipulate that in this new language, ‘existence’ strictly refers to **existence**. (Sider suggests that we

may also use ‘ \exists ’ to refer to **existence**, to avoid confusion.) We may then recast our ontological debates as being about **existence**, without getting distracted by what ‘existence’ means in ordinary English. According to Sider, there is therefore “a single best quantifier meaning”, which ‘existence’ should have, at least in the ontology room.

The central question of metaontology is that of whether there are many equally good quantifier meanings, or whether there is a single best quantifier meaning. It is a question about nature’s joints; it is a question of *how much quantificational structure the world contains*.

(Sider 2009, p. 397, emphasis in original).

Because the inconsistency approach is primarily a descriptive thesis, this normative proposal is not actually in conflict with the inconsistency approach. If the ordinary concept of existence is inconsistent, it would in fact seem rather unwise to keep using it inside the ontology room. Instead we need to engineer some new existence-like concept or concepts for those properties or structures that we would like to talk about, and Sider’s proposal thus seems to be on the right track.

Nonetheless, it is indeed still an open question whether there really only is one existence-like natural property, rather than several, or whether there is any such property or structure at all. The inconsistency approach leaves all of these options open. I would for this reason disagree with Sider that there necessarily has to be a “single best quantifier meaning”, i.e. **existence**, as there could just as well be several equally good, and equally substantive, quantifier meanings. Stipulating that there is a “single best quantifier meaning” therefore seems to prejudge the question that we are trying to answer. In addition, we of course also need to be asking, “best for what?” The answer would seem to depend on what hypothetical existence-like property we are interested in investigating. All ontologists may not necessarily be interested in the same property. Or perhaps it turns out that there is no such property or structure at all, or only non-natural existence-like properties. Significantly though, even if it turns out that there is no such existence-like natural property or structure at all, that in itself is a substantive ontological result which we presumably need to do some ontology in order to establish. In any case, there is at least *one* substantial and non-trivial ontological question to be answered.

A further objection to the view, which I mentioned very briefly in relation to quantifier variantism in chapter one, might be that existence cannot be an inconsistent concept because the existential quantifier has a perfectly consistent formal rule of use in logic. But this conflates the concept of existence with the existential quantifier (Quineans especially are guilty of this

conflation). The existential quantifier, ‘ \exists ’, isn’t strictly speaking the same as the concept of existence in ordinary English -- in fact, it isn’t even a concept, but a *model* of one. My claim is not that ‘ \exists ’ is inconsistent -- my claim is that EXISTENCE, as we use the concept in ordinary English, is inconsistent. In contrast to the concept of existence in ordinary English, ‘ \exists ’ is a philosopher’s invention. It has a fixed, formal rule of use in predicate logic, and that use is not inconsistent. We may of course insist that ‘ \exists ’ models our concept of existence -- but precisely because it is a philosopher’s invention, there is no independent guarantee that it is an adequate model. As should be clear by now, I in fact do not think that ‘ \exists ’ is an adequate model of our ordinary concept of existence at all, seeing as EXISTENCE is inconsistent, while ‘ \exists ’ is not.

Further reasons for thinking that it is not an adequate model is that there appear to be sentences in natural language involving our ordinary concept of existence which we cannot formalise in the standard way without something being lost along the way (Crane 2012; 2013). In fact, there seem to be natural language sentences -- such as negative existentials -- that are straightforwardly true, but which are incompatible with standard approaches to existence in formal semantics. Take for example the following sentence (the example is borrowed from Crane 2012):

(S): Some characters in *War and Peace* existed and some did not.

Sentence S, in ordinary English, appears to be a true sentence. The problem is that if we treat both ‘some’ and ‘exists’ as quantificational expressions, as the standard Quinean view recommends, we get a contradiction if we attempt to formalise sentences such as S. To be more precise, the standard view says that we should treat ‘Some Fs are Gs’ as another way of saying ‘There exists Fs which are G’. But if that is correct, then sentence S actually says the following:

(S*): There exists characters in *War and Peace* which exist and there exists characters in *War and Peace* which do not exist.

But sentence S* clearly contains a contradiction, and therefore sentence S* cannot be true. So it seems that sentence S* cannot be equivalent to sentence S in ordinary English, because in ordinary English, S does not seem to be saying anything contradictory. In fact, it seems to be saying something that is straightforwardly true. In ordinary English, then, it seems that there must be a distinction between quantification and existence, since in ordinary English we can quantify over things which do not exist without uttering contradictions (see Crane 2013 for an extended discussion of this point).

Of course, this does not necessarily mean that I think we should come up with a new model; in this case, it seems that because the model is consistent, the model is actually better, and more useful, than our ordinary inconsistent concept of existence. The point is merely that if we are engaged in a descriptive project of capturing our concept of existence, rather than a revisionist project, then ‘ \exists ’ is not an adequate model of existence.

A related worry would also have it that it is problematic for me to employ expressions such as ‘there is’ or ‘there are’ in the course of this discussion, given that EXISTENCE is inconsistent. My response to this worry is along the same lines as above -- ‘there is’ and ‘there are’ are not strictly the same as the concept of existence. In contrast to our ordinary concept of existence, the expressions ‘there is/are’ seem to me, at least *prima facie*, unproblematic and potentially more analogous to the formal use of the existential quantifier than the concept of existence (hence ‘ \exists ’ might be a better model of expressions such as ‘there is/are’ than of existence). Going into exactly how these expressions work, or exactly what their meaning is, would take us too far afield, but evidence from linguistics also seem to support the view that the semantics of ‘exists’ and ‘there is/are’ are fundamentally different (see e.g. Moltmann 2013).¹⁷

Now, one might argue that because philosophers standardly formalise expressions such as ‘there is’ in the same way as ‘exists’, this is evidence for them being synonymous. But this is merely a philosophical convention, a convention which may as well have come about because it is a convenient simplification, and not necessarily because it is the most adequate model for the semantics of existence (in addition, if the concept of existence is inconsistent, a completely adequate model, faithful to the inconsistency, is probably not going to be particularly useful). Therefore, I think it is far from obvious that the default position, at least with respect to ordinary English, should be that these expressions are unproblematically interchangeable. Moreover, if the expressions were obviously synonymous, the Meinongian position (or any position which makes a distinction between being and existence or quantification and existence) would just be a non-starter, because then Quine would obviously be correct. One need not be a Meinongian to

¹⁷ Further, Moltmann (2013) notes that: “In philosophy, there are two opposing views on existence. On one view, existence is a univocal concept and closely tied to existential quantification and counting. If there is one thing and there is another thing, even of a very different kind, then there are two things. On the other view, things of different kinds may ‘exist’ in different ways or engage in different ‘modes of being’. Whereas the former view is clearly the dominant one in contemporary analytic philosophy, various versions of the latter view can be found throughout the history of philosophy. Natural language, it appears, reflects both views, in two different types of sentences. *There*-sentences as well as sentences with simple existential quantifiers (*some, a*) may be used so as to reflect the first notion of existence; existence statements, that is, subject-predicate sentences with *exist* or *occur*, reflect the second notion.” (p. 33). This observation, that natural language reflects both the Quinean and the “Meinongian” views, is entirely in line with what is expected on the inconsistency approach.

acknowledge this¹⁸ -- but merely recognising that it isn't an *obvious* truth to say that existence is purely a quantificational expression acknowledges that in ordinary English there is room for a conceptual difference between quantification and existence -- however subtle that difference may be, or however convenient it might be to ignore the distinction for one's philosophical purposes.

A residual worry is that, even if we grant that the concept of existence is distinct from 'there is/are', if the concept of existence is inconsistent, the problems with 'exists' will carry over to expressions such as 'there is/are', given that we generally accept inferences between existence and these expressions. For example, we generally accept inferences like 'If x exists, then there is an x' and 'If x does not exist, then there is no x'. I think this is a genuine worry, but I doubt this poses a problem for my argument, given that I have at no point in my argument relied on such an inference. If EXISTENCE is indeed inconsistent, that also gives us reason to question the correctness of these inferences, and perhaps even reject them -- at least in the ontology room.

It can of course still be objected that it is hasty to assume that the problem has to lie with the concept of existence per se, instead of with e.g. the expression 'there is/are' or the concepts of fictional or imaginary entities. Nothing I have said so far strictly rules this out (although it appears to me that 'there is/are' in isolation are relatively unproblematic expressions). It might even seem plausible to argue that the problems arise from the interaction between expressions such as 'there is' with the concept of existence. Perhaps then it is rather our *conceptual scheme* that is inconsistent, in place of the concept of existence? I do not want to preclude this from being the case, but currently (as far as I'm aware) there is no theory of inconsistent conceptual schemes -- thus I would not really know how to further articulate this idea. What I do have is a theory of inconsistent concepts, and an argument pertaining to the concept of existence, which gives us reason to entertain the idea that EXISTENCE specifically is inconsistent. But presumably the same metaontological explanatory benefits would be available if the inconsistency turned out to occur at the level of conceptual schemes instead of at the level of individual concepts.

As a final challenge; if our concept of existence really is inconsistent, how come we haven't realised this sooner? The natural response to this question is that I think that we *have* realised that it is inconsistent. Philosophers have been puzzling over the paradox presented by the problem of non-being all along. The primary difference is that we have, at least to some extent, misdiagnosed this as a puzzle about existence in itself, instead of squarely viewing it as a

¹⁸ In fact, it is a feature of the inconsistency approach that one can separate expressions such as 'there is/are' from 'there exists' without agreeing with either Meinong or Quine.

problem rooted in our *concept* of existence. The tendency of some metaphysicians to turn to notions of fundamentality, grounding, or of “what is real” instead of discussing existence, is, I think, symptomatic of the fact that many others share the intuition that existence is somehow not a useful, or even a problematic, concept. In addition, those who have analysed the concept of existence have generally tried to come up with consistent analyses, which from the outset excludes the possibility that the concept of existence is inconsistent from even being considered. And ultimately, we could ask about *any* particular conceptual analysis of existence why we hadn’t thought of *that* analysis sooner -- but surely that shouldn’t be a reason to discredit the analysis!

3.5 Conclusion

In this chapter I have tried to show how the inconsistency approach to existence can help us make sense of ontological debates. The finer details of the precise inconsistency, and to what extent this is a feature isolated in the concept of existence or also present in other quantificational expressions, still remain to be worked out. But given the explanatory advantages over the other metaontological views which I have discussed, I believe the proposal that EXISTENCE is an inconsistent concept should be worth taking seriously.

CONCLUDING REMARKS

In this thesis I have argued that understanding our ordinary concept of existence as an inconsistent concept helps us to make better sense of what is going on in ontological debates. Not only does it enable us to tell a more charitable story about what ontologists are up to, but it also allows us to explain why the intuition that ontological debates are merely verbal, metalinguistic, or insubstantive is so widespread -- namely, because it is frequently correct. That notwithstanding, if the cause of the defectiveness of ontological debates can be attributed to the concept of existence itself, this is in part good news, as a straightforward solution immediately suggests itself; if we can fix the concept of existence, we can fix ontological debates.

Still, much work remains to be done in working out the precise nature of the inconsistency. Although I offered a crude example of an inconsistent conceptual analysis in chapter three, my aim in this thesis has primarily been to try and illustrate that the hypothesis that 'existence' expresses an inconsistent concept has explanatory advantages over other metaontological accounts. The overall structure of my argument presented here is an inference to the best explanation; the precise details of how the ordinary concept of existence functions remain to be worked out. In order to establish the exact nature of the inconsistency or the conceptual defect, I suggested earlier that we need to do naturalised conceptual analysis. Establishing precisely what the culprit is, is presumably going to be useful for figuring out how to best fix the defect.

Much of the difficult work, however, is likely going to be in the form of actually devising functional, non-defective conceptual revisions, or even replacements. Although these revisions or replacements are likely to be superfluous in ordinary contexts and conversations, it might indeed be wise for philosophers to switch into some form of Ontologese inside the philosophy room. To some extent, I believe this is already happening, for instance in the form of Sider's proposal to focus on joint-carving terms, and the general shift in focus within metaphysics away from questions of existence, to questions of fundamentality, grounding, and reality.

Unfortunately, I have no elaborate replacement or revision proposals to offer here, but one suggestion would be to frame ontological questions in terms of domains. A natural way of thinking of existence and existence-like concepts is to think of them as designating a domain (a little bit like Chalmers' furnishing functions). Instead of asking whether numbers, ordinary objects, and universals are part of some inconsistently defined domain of "what exists", we could then ask whether they are part of some precisely and consistently defined domain that we are interested in, such as the domain of that which is fundamental, the domain of that which is physical or mind-independent, or the domain of things which ground all other things (although, this of course presupposes that we first make explicit precisely what we mean by 'fundamentality', 'physicality', 'mind-independence' or 'ground'). Further, we could then also ask if any of these domains are in some sense more "absolute" than any of the others (provided that we first define what we mean by 'absolute').

There is thus much work left to be done in figuring out how to best replace or revise our ordinary concept of existence, and what we want these new concepts to do for us. Fortunately for the ontologists, there is currently a lot of new research on the topic of conceptual engineering and conceptual ethics (see e.g. Burgess & Plunkett 2013a; 2013b; Simion 2017; Cappelen 2018; Burgess, Cappelen & Burgess forthcoming), which one could hopefully draw on in engineering new concepts for doing ontology.

A worry which one may have about the proposed engineering project is that without the ordinary concept of existence, it will no longer be "ontology" which we are doing. That is, the worry is that by changing the concept, we are changing the topic. This is known in the literature as the *discontinuity objection* (see e.g. Prinzing 2017). The worry can also be expressed as the worry that this new form of ontology, with new concepts, would fail to answer the original ontological questions. However, if our ordinary concept of existence is a defective concept, I think it is quite clear that this worry is unfounded. There is no need to worry about not answering the original questions, as the original questions are defective (or pseudo-questions, if you like); they have no real answers (perhaps this was part of Chalmers' motivation for embracing anti-realism).

Nonetheless, I think that most of those who have the intuition that ontological questions are real and substantive are correct insofar as there is indeed a subject matter here that is real and substantive. Or, at least, none of the deflationary arguments considered in this thesis have been successful in showing that there is no such subject matter. Moreover, in thinking about ontological matters, the inconsistency approach predicts that most ontologists have actually

already formed their own consistent conceptions of existence, and most likely succeeded in asking substantive ontological questions by using these consistent conceptions -- even if they are failing to properly express these questions in a common language. The crucial next step would be to precisely articulate these consistent conceptions into technical concepts, which can be used for the purposes of doing ontology. Nothing therefore seems to rule out that there are substantive ontological questions in the vicinity to be asked and answered -- even if these don't employ our ordinary concept of existence. (Whether this quasi-new subject matter may rightfully be called 'ontology' is a verbal dispute which I don't see much point in pursuing.)

Nevertheless, worth emphasising still is that just because these questions do not necessarily turn out to be pseudo-questions, that does not give them a free pass. Their relative importance to other philosophical questions still of course needs to be assessed, and nothing much about their significance or depth actually follows from the fact that they aren't necessarily trivial pseudo-questions (surely that's a *very* low bar for any inquiry).

Another important remark is that these conclusions about ontological questions apply mostly to traditional ontological debates about the existence of numbers, holes, ordinary objects, universals, and so on. Social ontology, for instance, would appear to be largely unaffected, since their project was never to establish what "exists" in some absolute sense in the first place, but rather to answer normative questions about how we ought to think of entities and phenomena that are part of our social reality.

Insofar as the wider metaphilosophical aspirations of this thesis are concerned, I wish to leave it up to the reader to pass judgment. What should be clear from my discussion in chapter two, however, is that if one believes that focusing on conceptual defects may help us make progress on other philosophical topics, there is also much work to be done on the nature of concepts, conceptual defects, and inconsistent concepts. We need more developed theories to form a more nuanced picture of the range of possible conceptual defects, and to what extent we can make use of inconsistency approaches to make sense of other philosophical debates.

Finally, a general lesson to be drawn is that it would most likely serve us well to be more self-conscious of the concepts that we employ. By taking a critical stance towards our conceptual repertoire, and more critically evaluating and engineering concepts, it seems that we should be able to improve philosophical debates -- or, minimally, protect ourselves from the accusation (or realisation!) that our debates are "merely verbal", or about some indeterminate subject matter.¹⁹

¹⁹ Thanks to my supervisor, Daniel, for valuable discussion and comments on all parts of this thesis.

BIBLIOGRAPHY

- Azzouni, J. (2007). 'The Inconsistency of Natural Languages: How We Live With It', *Inquiry* 50(6), pp. 590-605.
- Balcerak Jackson, B. (2014). 'Verbal Disputes and Substantiveness,' *Erkenntnis* 79, pp. 31-54.
- Belleri, D. (2017). 'Verbalism and Metalinguistic Negotiation in Metaontological Disputes', *Philosophical Studies* 174(9), pp. 2211-2226.
- Bennett, K. (2009). 'Composition, Colocation, and Metaontology,' in D. Chalmers, D. Manley and R. Wasserman (eds.), *Metametaphysics: New Essays on the Foundations of Ontology*, pp. 38-76. Oxford University Press.
- Blatti, S., & Lapointe, S. (eds.). (2016). *Ontology After Carnap*. Oxford University Press.
- Burge, T. (1993). 'Content Preservation,' *The Philosophical Review* 102(4), pp. 457-488.
- Burgess, A., Cappelen, H. & Plunkett, D. (eds.). (forthcoming). *Conceptual Ethics and Conceptual Engineering*. Oxford University Press.
- Burgess, A. & Plunkett, D. (2013a). 'Conceptual Ethics I', *Philosophy Compass* 8(12), pp. 1091-1101.
- (2013b). 'Conceptual Ethics II', *Philosophy Compass* 8(12), pp. 1102-1110.
- Cappelen, H. (2018). *Fixing Language: An Essay on Conceptual Engineering*. Oxford University Press.
- Carnap, R. (1950). 'Empiricism, Semantics, and Ontology', *Revue Internationale de Philosophie* 4(11), pp. 20-40.
- Chalmers, D., Manley, D., & Wasserman, R. (eds.). (2009). *Metametaphysics: New Essays on the Foundations of Ontology*. Oxford University Press.
- Chalmers, D. (2009). 'Ontological Anti-Realism,' in D. Chalmers, D. Manley and R. Wasserman (eds.), *Metametaphysics: New Essays on the Foundations of Ontology*, pp. 77-129. Oxford University Press.
- (2011). 'Verbal Disputes', *The Philosophical Review* 120(4), pp. 515-566.

- Crane, T. (2012). 'Existence and Quantification Reconsidered,' in T. Tahko (ed.), *Aristotelian Metaphysics*, pp. 44-65. Cambridge University Press.
- (2013). *The Objects of Thought*. Oxford University Press.
- Dorr, C. (2005). 'What We Disagree About When We Disagree About Ontology,' in M. Kalderon (ed.), *Fictionalism in Metaphysics*, pp. 234-286. Oxford University Press.
- Eklund, M. (2002). 'Inconsistent Languages,' *Philosophy and Phenomenological Research* 64(2), pp. 251-275.
- (2009). 'Carnap and Ontological Pluralism,' in D. Chalmers, D. Manley, & R. Wasserman, *Metametaphysics: New Essays on the Foundations of Ontology*, pp. 130-156. Oxford University Press.
- Fodor, J. A. (1975). *The Language of Thought*. Cambridge, MA: Harvard University Press.
- (1998). *Concepts: Where Cognitive Science Went Wrong*. Oxford University Press.
- Greenough, P. (2017). 'Conceptual Marxism and Truth: Inquiry Symposium on Kevin Scharp's Replacing Truth,' *Inquiry*, DOI: 10.1080/0020174X.2017.1287919.
- (forthcoming). 'Neutralism and Conceptual Engineering,' in A. Burgess, H. Cappelen, and D. Plunkett (eds.), *Conceptual Ethics and Conceptual Engineering*. Oxford University Press.
- Hirsch, E. (2002a). 'Quantifier Variance and Realism,' *Philosophical Issues* 12, pp. 51-73.
- (2002b). 'Against Revisionary Ontology,' *Philosophical Topics* 30(1), pp. 103-127.
- (2005). 'Physical-Object Ontology, Verbal Disputes, and Common Sense,' *Philosophy and Phenomenological Research* 70(1), pp. 67-97.
- (2008). 'Language, Ontology, and Structure,' *Noûs* 42(3), pp. 509-528.
- (2009). 'Ontology and Alternative Languages,' in D. Chalmers, D. Manley, & R. Wasserman, *Metametaphysics: New Essays on the Foundations of Ontology*, pp. 231-259. Oxford University Press.
- (2011). *Quantifier Variance and Realism: Essays in Metaontology*. Oxford University Press.
- Horden, J. (2014). 'Ontology in Plain English,' *Philosophical Quarterly* 64, pp. 255-242.
- Lewis, D. K. (1986). *On the Plurality of Worlds*. Blackwell Publishing.
- McDaniel, K. (2009). 'Ways of Being,' in D. Chalmers, D. Manley, & R. Wasserman, *Metametaphysics: New Essays on the Foundations of Ontology*, pp. 290-319. Oxford University Press.
- (2010a). 'A Return to the Analogy of Being,' *Philosophy and Phenomenological Research* 81(3), pp. 688-717.

- (2010b). 'Being and Almost Nothingness', *Noûs* 44(4), pp. 628-649.
- Machery, E. (2009). *Doing Without Concepts*. Oxford University Press.
- (2017). *Philosophy Within Its Proper Bounds*. Oxford University Press.
- Meinong, A. (1904/1960). 'The Theory of Objects (I. Levi, D. B. Terrell & R. M. Chisholm, trans.)', in R. Chisholm (ed.), *Realism and the Background of Phenomenology*, pp. 76-117. Free Press.
- Merricks, T. (2001). *Objects and Persons*. Oxford: Clarendon Press.
- Moltmann, F. (2013). 'The Semantics of Existence', *Linguistics and Philosophy* 36(1), pp. 31-63.
- Parfit, D. (1987). 'Divided Minds and the Nature of Persons', in C. Blakemore & S. A. Greenfield (eds.), *Mindwaves*, pp. 19-26. Blackwell Publishing.
- Parsons, T. (1975). 'A Meinongian Analysis of Fictional Objects', *Grazer Philosophische Studien* 1, pp. 73-86.
- (1980). *Nonexistent Objects*. New Haven: Yale University Press.
- (1982) 'Are There Nonexistent Objects?' *American Philosophical Quarterly* 19, pp. 365-371.
- Patterson, D. (2009). 'Inconsistency Theories of Semantic Paradox', *Philosophy and Phenomenological Research* 79(2), pp. 387-422.
- Plato (1993). *Sophist* (Trans. N. P. White). Hackett Publishing.
- Plunkett, D. (2015). 'Which Concepts Should We Use?: Metalinguistic Negotiations and The Methodology of Philosophy', *Inquiry* 58, pp. 828-874.
- Plunkett, D. & Sundell, T. (2013). 'Disagreement and the Semantics of Normative and Evaluative Terms', *Philosophers' Imprint* 13(23), pp. 1-37.
- Prinz, J. (2002). *Furnishing the Mind: Concepts and Their Perceptual Basis*. Cambridge, MA: The MIT Press.
- Prinz, M. (2017). 'The Revisionist's Rubric: Conceptual Engineering and the Discontinuity Objection', *Inquiry*. DOI: 10.1080/0020174X.2017.1385522
- Quine, W. V. (1948). 'On What There Is,' *The Review of Metaphysics* 2(5), pp. 21-38.
- Scharp, K. (2007). 'Replacing Truth', *Inquiry* 50(6), pp. 606-621.
- (2013). *Replacing Truth*. Oxford University Press.

- (forthcoming). 'Philosophy as the Study of Defective Concepts,' in A. Burgess, H. Cappelen, and D. Plunkett (eds.), *Conceptual Engineering and Conceptual Ethics*. Oxford University Press.
- Sider, T. (2009). 'Ontological Realism,' in D. Chalmers, D. Manley, & R. Wasserman, *Metametaphysics: New Essays on the Foundations of Ontology*, pp. 384-423. Oxford University Press.
- (2011). *Writing the Book of the World*. Oxford University Press.
- Simion, M. (2017). 'The 'Should' in Conceptual Engineering', *Inquiry*, DOI: 10.1080/0020174X.2017.1392894
- Soames, S. (1999). *Understanding Truth*. Oxford University Press.
- Sorensen, R. A. (1992). *Thought Experiments*. Oxford University Press.
- Thomasson, A. L. (2007). *Ordinary Objects*. Oxford University Press.
- (2009). 'Answerable and Unanswerable Questions', in D. Chalmers, D. Manley, & R. Wasserman, *Metametaphysics: New Essays on the Foundations of Ontology*, pp. 444-471. Oxford University Press.
- (2015). *Ontology Made Easy*. Oxford University Press.
- (2017). 'Metaphysical Disputes and Metalinguistic Negotiation', *Analytic Philosophy* 58(1), pp. 1-28.
- van Inwagen, P. (1998). 'Meta-Ontology,' *Erkenntnis* 48, pp. 233-250.
- (2009). 'Being, Existence, and Ontological Commitment', in D. Chalmers, D. Manley, & R. Wasserman, *Metametaphysics: New Essays on the Foundations of Ontology*, pp. 472-506. Oxford University Press.
- Williamson, T. (2007). *The Philosophy of Philosophy*. Oxford University Press.