Ontological Anti-Realism*

David J. Chalmers

1 Introduction

The basic question of ontology is "What exists?". The basic question of metaontology is: are there objective answers to the basic question of ontology? Here ontological realists say yes, and ontological anti-realists say no.

(Compare: The basic question of ethics is "What is right?". The basic question of metaethics is: are there objective answers to the basic question of ethics? Here moral realists say yes, and moral anti-realists say no.)

For example, the ontologist may ask: Do numbers exist? The Platonist says yes, and the nominalist says no. The metaontologist may ask: is there an objective fact of the matter about whether numbers exist? The ontological realist says yes, and the ontological anti-realist says no.

Likewise, the ontologist may ask: Given two distinct entities, when does a mereological sum of those entities exist? The universalist says always, while the nihilist says never. The metaontologist may ask: is there an objective fact of the matter about whether the mereological sum of two distinct entities exists? The ontological realist says yes, and the ontological anti-realist says no.

Ontological realism is often traced to Quine (1948), who held that we can determine what exists by seeing which entities are endorsed by our best scientific theory of the world. In recent years, the practice of ontology has often presupposed an ever-stronger ontological realism, and strong versions of ontological realism have received explicit statements by Fine (this volume), Sider (2001; this volume), van Inwagen (1998; this volume), and others.

Ontological anti-realism is often traced to Carnap (1950), who held that there are many different ontological frameworks, holding that different sorts of entities exist, and that while some

[0]INPCers: For a 30-page version, the most important sections to read are probably 1, 2, and 5-10. For those who have read the version of this paper posted in early January, the current version (early February) has only minor changes from that version.

frameworks may be more useful than others for some purposes, there is no fact of the matter as to which framework is correct. In recent years, versions of ontological anti-realism have been developed by Putnam (1987), Sidelle (2001), Yablo (this volume), and others.

An intermediate sort of *lightweight realism* has also developed, holding that while there are objective answers to ontological questions, these answers are somehow shallow or trivial, perhaps reflecting conceptual truths rather than the furniture of the world. Deflationary views of this sort have been developed by Hirsch (1993; this volume), Thomasson (this volume), Wright and Hale (2001; this volume), and others. These views contrast with what we might call the *heavyweight realism* of Fine, Sider, van Inwagen, and others, according to which answers to ontological questions are highly nontrivial, and reflect the ultimate furniture of the world.

In the currently thriving field of first-order ontology, the most popular view is heavyweight realism, with a minority of lightweight realists and anti-realists. Outside the field of ontology, deflationary views are widespread, with many non-ontologists being skeptical of the heavyweight realism that has become common in the field. It is natural to suppose that there is some sort of selection effect at work here: those who think that ontological questions are deep questions with determinate answers are much more likely to go into ontology than those who think that the questions are shallow or lack objective answers.

Each of the views has its advantages and disadvantages. I will briefly mention reasons for and against each.

An intuition favoring deflationary views against heavyweight realism is the following. Say that we know all about the qualitative properties of two objects—two cups, say—and the qualitative relations between them, leaving out any properties or relations concerning objects that they jointly compose. There is a strong intuition that we are thereby in a position to know *everything* relevant there is to know about the objects. There is no deep further truth concerning whether the objects compose a further object (a cupcup, say) of which we are potentially ignorant. The question of whether there is a cupcup is a matter for bookkeeping or for semantic decision, perhaps, but it is not a matter for discovery.

An intuition favoring realist views against ontological anti-realism is the following. Some ontological theses can be stated in wholly logical vocabulary, with no problematic terms, or at least entail such theses. For example, the thesis of universalism entails: $\forall x \forall y \exists z (z = sum(x, y))$. Here the notion of a sum can be defined in terms of the basic notion of a part, so that if the notion of a part is a logical notion, this is a logical thesis. If one does not count the notion of part as logical, one can still note that ontological theses include wholly logical theses such as $\exists x \exists y (x \neq y)$.

Ontological anti-realism, at least in its strong form, is committed to denying that these theses have objective and determinate truth-values. But these theses have wholly logical vocabulary, in which no term is vague or context-dependent, and theses of this sort cannot fail to have an objective and determinate truth-value.

An intuition favoring heavyweight realism and anti-realism against lightweight realism is the following. Consider ontological theses such as 'Numbers exist', or 'If there are particles arranged as a heap, there is a heap', or 'If there are two objects, there is a sum of those two objects'. Lightweight realists hold that theses of this sort are trivially true (or false), and often hold that they are analytically true (or false). But no existence claim is trivially true, and certainly no existence claim is analytically true. It may be analytic or trivial that 'If there is an object with certain properties (numerical properties, shaped as a heap, composed wholly of those two objects), then there is a number/heap/sum.' But the claim that there is an object with those properties is never trivial or analytic, and is never trivially or analytically entailed by a sentence that does not make a corresponding existence claim.

My own sympathies are with ontological anti-realism, and especially with ontological anti-realism of a Carnapian variety. In the first half of this paper, I will do some logical geography, examining the language used to state ontological claims, and differentiating more carefully between what I have called heavyweight realism, lightweight realism, and ontological anti-realism. I hope that much of this material should be acceptable to theorists of quite different stripes. In the second half of the paper, I will flesh out ontological anti-realism and defend it against arguments against it, including the arguments above. I will also sketch some details of the semantics and metaphysics of an ontological anti-realist views.

In this paper, I will not *argue* for ontological anti-realism at any length. I will occasionally point to considerations in its favor, but detailed arguments for the view are the subject for a separate analysis. Instead, I will be more concerned to elaborate the view, making the case that it is internally coherent and that there are no strong reasons to reject it. I will take for granted that there are at least some strong prima facie attractions to ontological anti-realism, but that there are also serious challenges to whether the view is coherent and defensible, and I will mostly focus on answering the challenges. I will give some reasons for favoring ontological anti-realism over its close cousin, lightweight realism, but a full-scale argument against heavyweight realism is a matter for further work.

2 Ontological and Ordinary Existence Assertions

I will start with Carnap's distinction between internal and external questions, from "Empiricism, Semantics, and Ontology". On Carnap's view, questions about existence always involve linguistic frameworks: for example, the framework of mathematics, the framework of propositions, or the framework of commonsense objects. There are then two sorts of existence questions.

Internal questions are those posed within a framework, concerning the existence of certain specific entities within the framework. Examples include 'Are there any odd perfect numbers?' asked by a mathematician, and 'Is there an apple on the table?' asked by a child. Internal claims are answers to internal questions. On Carnap's view, internal claims are typically true or false. In some cases, such as mathematics, they will be analytically true or false, with their truth-value determined wholly by the rules of the framework. In other cases, such as claims concerning ordinary objects, they will be empirically true or false, with their truth or falsity determined by the rules of the framework in conjunction with experience and perhaps with other aspects of the world.

External questions are those posed outside a given framework, concern the existence of the framework's system of entities as a whole. Examples include 'Do numbers exist?', or 'Do ordinary physical objects exist?', asked from a purported neutral perspective. External claims are answers to external questions. On Carnap's view, external claims are neither true nor false. For Carnap, the choice between frameworks is practical rather than factual. Any purported factual question about which framework is the "correct" one is held to be a pseudoquestion, without cognitive content.

Even if one does not accept Carnap's claims about the properties of internal and external questions, there is something natural about the distinction itself. The distinction between internal and external questions seems to reflect a distinction in our practice of raising questions about existence, if nothing else. At the same time, Carnap's terminology is suboptimal. For a start, the terminology of "internal" and "external" is too closely tied to Carnap's theoretical apparatus involving frameworks to serve as a neutral starting point. If one rejects the idea of a framework, or of a question being internal to a framework, one will reject this version of the distinction. If possible, it is desirable to draw a relatively pretheoretical version of the distinction that almost anyone can accept, regardless of their theoretical inclinations.

In addition, "internal question" and "external question" tends to suggest two different sorts of *sentence*, whereas I think the most important distinction is between different *uses* of sentence (or perhaps, between different *evaluations* of sentences). For example, a sentence such as "Do prime

numbers exist?" might be used to pose an internal question (by a mathematician, say) or to pose an external question (by a metaphysician, say).[1] The same goes in principle for sentences such as 'Numbers exist' and 'There are four prime numbers less than ten.'

Instead, I will distinguish between two different sorts of existence assertions: *ordinary* and *ontological* existence assertions.[2] Here an assertion is an utterance of an assertive sentence. An existence assertion is an utterance of a sentence that appears to assert or deny the existence of certain entities: for example, 'Xs exist', 'There are Ys', 'There are no Zs'. One can also straightforwardly extend the distinction to sentences that appear to involve universal quantification, and to sentences in which apparently quantified claims are embedded. For the purposes of this definition, sentences are individuated by surface structure. So it is compatible with the definition that ontological and ordinary assertions of the same sentence may have the same surface structure while differing in deep structure (perhaps involving a difference in covert variables, operators, and the like).

An *ordinary* existence assertion, to a first approximation, is an existence assertion of the sort typically made in ordinary first-order discussion of the relevant subject matter. For example, a typical mathematician's assertion of 'There are four prime numbers less than ten' is an ordinary existence assertion, as is a typical drinker's assertion of 'There are three glasses on the table'.

An *ontological* existence assertion, to a first approximation, is an existence assertion of the sort typically made in broadly philosophical discussion where ontological considerations are paramount. For example, a typical philosophers' assertion of 'Abstract objects exist' is an ontological existence assertion, as is a typical philosophers' assertion of 'For every set of objects, there exists an object that is their mereological sum'.

We can think of ontological existence assertions as those made inside the "ontology room", and ordinary existence assertion as those made outside the ontology room. At the very least, there is a clear pragmatic difference between these two sorts of assertion. For example, given an ontological assertion of 'There are infinitely many prime numbers', it is appropriate to respond 'No, there aren't, because if numbers exist, they are abstract objects, and there are no abstract objects'. But given an ordinary assertion of 'There are infinitely many prime numbers', it is not

[1] It is arguable that Carnap himself saw the distinction this way. For example, he speaks of meaning 'There are numbers' in an internal sense or an external sense. This point is missed by Quine in "Carnap's Views on Ontology" (1951), who interprets Carnap's distinction as a syntactic distinction between sentences with different surface structures.

[2] Distinctions closely related to the distinction between ordinary and ontological existence assertions have been made by any number of philosophers. Apart from Carnap, see for example Dorr 2005, Hofweber 2005, Horgan 2001, Yablo 2001, and many others.

appropriate to respond in this way.

Correspondingly, it is natural to hold that ordinary and ontological existence assertions differ with respect to an important sort of utterance evaluation, which I will call *correctness*. The correctness of ontological existence assertions is sensitive to ontological matters, and indeed is obviously sensitive in this way. The correctness of ordinary existence assertions is insensitive to ontological matters, or at least is not obviously so sensitive.

For example, the correctness of an ordinary assertion of 'There are prime numbers' is insensitive to whether Platonism or nominalism is true. Even if nominalism is true, so that strictly speaking there are no numbers, an ordinary mathematician's assertion of 'There are infinitely many prime numbers' is correct. (Note that correctness need not be the same thing as truth.) Likewise, the correctness of an ordinary assertion of 'There are three glasses on the table' is insensitive to the truth or falsity of nihilism. Even if nihilism is true, so that strictly speaking there are no macroscopic objects, an ordinary drinker's assertion of 'There are three glasses on the table' may be correct.

By contrast, the correctness of an ontological assertion of 'There are prime numbers' is sensitive to whether Platonism or nominalism is true. If nominalism is true, a metaphysician's utterance of this sentence is incorrect. Likewise, the correctness of an ontological assertion of 'There are three glasses on the table' is sensitive to whether nihilism is true or false. If nihilism is true, then a metaphysician's assertion of this sentence is incorrect.

Of course I have not said just what correctness is, and I have not said how it is related to truth. It is highly plausible that an ontological existence assertion is correct if and only if it is true, but this is not so obvious for ordinary existence assertions. One might hold that here, correctness is some sort of pragmatic evaluation that need not coincide with truth. For example, it might be taken to involve the truth of an implicated content, or of some other associated content. Or it might be taken to coincide with a sort of intuitive acceptability in light of relevant empirical truths and first-order reasoning. On views of this sort, then if nihilism is true, an ordinary assertion of 'There are three glasses on the table' is not true, but it is nevertheless correct.

I will remain neutral between these views here. I am inclined to accept that correctness coincides with truth, but little here will depend on that. Whether or not correctness is truth, the important point is that there is *some* natural sort of utterance evaluation that behaves as I have said correctness behaves. I take it that this much is intuitively obvious.[3] At the very least, the correctness of an assertion is truth-like in various respects. For example that correctness or incorrectness depends on the state of the world, and not just on what a speaker currently finds acceptable. For

example, if conversational participants are subject to an appropriate illusion, an ordinary assertion of 'There are three glasses on the table' when the table is empty may be acceptable, but it is not correct. Intuitively, ordinary assertions have correctness-conditions, corresponding to those states of the world under which they will be correct. These correctness-conditions function very much like truth-conditions in ordinary discourse, reflecting a highly salient way that speakers take the world to be when they make a sincere assertion, and also reflecting the worldly conditions that hearers require for the acceptance or rejection of an given assertion.

I will also remain neutral, at least for now, on the correct linguistic account of the difference between ontological and ordinary assertions. On one view, this is a matter of context-dependence: a sentence such as 'Prime numbers exists' has different contents in different contexts, depending (for example) on standards that are operative in those contexts. On another view, the difference arises from an ambiguity in 'exists' and in related expressions. On a third view, the difference arises from different pragmatic standards for acceptibility in different contexts, perhaps arising from a difference in implicated contents or other associated contents. On other views, the difference involves some other form of semantic or pragmatic underdetermination. Again, I am inclined to accept the first view, but little here will depend on that.

It is also possible to hold that the relevant difference is not fundamentally a difference between two sorts of *utterances*, but really a difference between two ways of *assessing* sentences and/or utterances. For example, one could hold that the sentence 'There are prime numbers' is correct when assessed by ordinary standards, but (if nihilism is true) incorrect when assessed by ontological standards. One can then hold that a mathematician's utterance of this sentence is correct by ordinary standards but not by ontological standards, and that the same goes for a philosopher's utterance. Here the relevant standards are those of a context of assessment of the utterance, which need not be the context of the utterance itself. If one holds that correctness is a sort of truth, this view leads to a sort of relativism about the truth of existence assertions, whereby they have different truth-values relative to different standards of assessment. Even on this view, though, there will be a difference between ontological and ordinary existence assertions, depending on whether the standards of assessment operative in the speaker's own context are ordinary standards or ontological standards. So I will continue to talk about two kinds of assertion, though I will return to this

[3]For views on which correctness is truth, and on which the truth-conditional contents of existence assertions vary between ordinary and ontological contexts, see Hofweber 2005 and Horgan 2001. For views on which correctness is something other than truth, and on which ordinary and ontological existence assertions have the same truth-conditions but different correctness-conditions, see Rayo forthcoming and Yablo 2001.

sort of view later.

Although the correct theoretical gloss on the distinction between ordinary and ontological existence assertions is disputable, the distinction itself is relatively pre-theoretical. The distinction should be acceptable to theorists of many different stripes. In particular, the distinction should be just as acceptable to ontological realists as to ontological anti-realists.

This is particularly clear in the case of those realists who endorse a revisionary metaphysics: roughly, a view on which the correct ontological theory denies some claims of commonsense ontology. These theorists usually need a sense in which ordinary assertions of a sentence S can be correct, even though an ontological assertion of S is false. For example, it is crucial for nominalists to allow a sense in which a mathematician's utterance of 'There are infinitely many prime numbers' is correct. It is likewise crucial for a nihilist to allow a sense in which 'There are two apples on the table' can be correct, and for a universalists to allow a sense in which 'There are two objects on the table' can be correct in the same circumstances. If one cannot do this, then one cannot account for a key feature of our ordinary use of these sentences.

Of course different revisionary metaphysicians may give different theoretical accounts of correctness. On one view, the correctness of an utterance may amount to the truth of a paraphrased sentence. (For example, a nihilist may hold that an ordinary utterance of 'There is a table' is correct iff 'There are particles arranged table-wise' is true.) On another view, it may amount to the truth of a conditional claim. (For example, a nominalist may hold that an ordinary utterance of 'There are prime numbers' is correct iff it is true according to the fiction of mathematics that there are prime numbers.) On a third view, it may amount to the truth of a content for which the domain of quantification is restricted. (For example, a universalist may hold that an ordinary utterance of 'There are two objects on the table' is true iff there are two familiar macroscopic objects on the table.) But all of these views in effect recognize a distinction between ordinary and ontological existence assertions. All allow that the same sentence can be uttered truly in ontological contexts and falsely, or at least incorrectly, in otherwise identical ordinary contexts.

The distinction is most likely to be questioned by realists who endorse a descriptive metaphysics: one on which correct ontological theory coincides with commonsense ontology (perhaps with differences stemming from empirical rather than philosophical sources, allowing the correct theory to accept photons and reject witches.) Strictly speaking, these realists can accept the distinction between assertions made inside and outside the ontology room. They will simply hold that the assertions have the same correctness conditions: the correctness conditions for each will mirror the commitments of commonsense ontology.

Here it is useful to distinguish two sorts of descriptive metaphysician. Some descriptive metaphysicians will hold that the coincidence between commonsense and correct ontology is a non-trivial fact about the world. On this view, one may reasonably hold that ontological and ordinary existence assertions differ in cognitive significance, though it turns out that their truth-values and correctness-values coincide. On this view, the correctness of ordinary assertions is trivially sensitive to the commitments of commonsense ontology, while the correctness of ontological assertions is nontrivially sensitive to commitments of commonsense ontology. So descriptive metaphysicians of this sort can clearly accept the distinction between ordinary and ontological existence assertions, and between different ways that these are associated with conditions of correctness.

Other descriptive metaphysicians will hold that the coincidence between commonsense and correct ontology is a trivial fact. On this view, the only sense that one can give to ontological existence assertions derives from that of ordinary existence assertions, so that the two sorts of claim coincide in cognitive significance, and automatically coincide in correctness. On this view, there may be no relevant difference in the conditions of correctness for ordinary and ontological existence assertions. This sort of view is an important view, but it clearly requires a fairly deflationary view of ontology that is closer to ontological anti-realism or lightweight realism than to heavyweight ontological realist. So I will set views of this sort aside for now, and return to them later.

3 Disagreements in Commonsense Ontology

We have seen that the correctness of an ordinary existence assertions is not trivially sensitive to the truth of ontological theories. Reflecting this fact, it is notable that in our community, proponents of very different ontological views typically agree about judgments about the correctness of ordinary existence assertions in specific circumstances. For example, Platonists and nominalists agree on the correctness of ordinary assertions (though not ontological assertions) of 'There are infinitely many primes'. Likewise, nihilists, universalists, and others can agree on the correctness of an ordinary assertion (though not an ontological assertion) of 'There are two objects on the table'.

We might explain this by noting that the correctness of an ordinary existence assertion coincides with its intuitive acceptibility in light of certain empirical truths about the world (for example, truths about the qualitative distribution of matter) and first-order reasoning (for example, mathematical reasoning), but not in light of any distinctively ontological truths and distinctively ontological reasoning. Let us say that our *commonsense ontology* is constituted by our dispositions

to accept ordinary existence claims in light of this sort of empirical information and first-order reasoning. In our community, commonsense ontology appears to be committed to particles, ordinary macroscopic objects such as tables, and at least some abstract objects such as numbers, but it does not appear to be committed to more recherche objects such as arbitrary mereological sums.

On this picture, the correctness of an ordinary existence assertion is constitutively connected to the commitments of commonsense ontology, but not to the commitments of ontological theory. For example, commonsense ontology is committed to glasses but not to mereological sums of glasses. So an ordinary assertion of 'There are two objects made of glass' is correct in a situation with two single glasses on the table, while an ordinary assertion of 'There are three objects made of glass' is not.

This constitutive connection raises an immediate question. It is plausible that different speakers and different communities might have different commonsense ontologies. Does it follow that the correctness of an ordinary existence assertion depends in some way on the speaker or the community? That is, could ordinary assertions of the same existential sentence by different speakers differ in their correctness, even when the nonquantificational elements of the sentence do not involve any context-dependence?

Let us suppose that for Martians but not humans, commonsense ontology includes arbitrary mereological sums. Faced with two apples on a bare table, and asked "How many objects are on the table", humans and Martians will usually make the following ordinary assertions (note that these are not ontological assertions):

> Human: There are exactly two objects on the table
>
> Martian: There are exactly three objects on the table.

We can then ask: which of these ordinary assertions is correct? Presumably the human's assertion is correct, but is the Martian's? Only two answers seem to be tenable. According to the first answer, both assertions are correct. On this view, the correctness of an ordinary existence assertion depends on the speaker's context and/or community. According to the second view, the human's assertion is correct, while the Martian's utterance is incorrect. Even on the second view, however, one should allow that the Martian's utterance is correct *by Martian standards*: that is, it is not h-correct, but it is m-correct. On this second view, there will be multiple notions of correctness, possessed by different evaluators.

I think that the first view is the most natural, but I will not choose between them here. On either view, there is a sort of variability in the standards associated with correctness. On the

first view, correctness of an ordinary existence assertion is tied to the commonsense ontology of the speaker or the speaker's community. On the second view, there is more than one notion of correctness (that is, more than one notion that might be expressed by terms such as 'correctness'), so that a human's notion is tied to human commonsense ontology, and the Martian's notion is tied to Martian commonsense ontology. Either way, the two assertions above will be on a par from a "God's eye" point of view, at least where standards in the vicinity of correctness are concerned.

One can then ask: Do the human and the Martian have a substantive disagreement? On the face of it, the mere fact that they make the ordinary existence assertions above does not entail this. If an ordinary human and Martian, making these assertions, are confronted with each other, they may well resolve the apparently disagreement terminologically. For example, they might well say something like:

> "It depends on how you count objects. The way you count objects, you're right; the way I count objects, I'm right. We could put things neutrally by saying that there are two h-objects and three m-objects. Or even better, we can say that there are three objects on the table by your standards, but not by my standards".

In this case, the appearance of a disagreement will fall away. Of course if the disputants are ontologically inclined, there might be some residual disagreement. But any such disagreement would likely involve disagreements over matters of ontological theory—for example, about whether entities that are m-objects but not h-objects really exist. This disagreement would lie in the realm of ontological existence assertions, and would not arise simply in virtue of the difference in ordinary existence assertions.

This suggests that the difference in ordinary existence assertions between humans and Martians is in a certain sense superficial. We have already seen that making an ordinary existence assertion does not commit a subject to a corresponding ontological existence assertion. Likewise, the apparent disagreement between humans and Martians need not involve a disagreement over ontological theory, and may well involve no substantive disagreement at all. The difference in commonsense ontology reflects a difference in classificatory practices rather than a difference in deep commitments.

In effect, the character of this disagreement is closer to the sort of disagreement one might find with a term such as 'tall' than the sort of disagreement one finds with a term such as 'wrong'. When a member of the academic community says 'John is tall' and a member of the basketball community says 'John is not tall', they are unlikely to take this to be a substantive disagreement.

Instead they are likely to resolve the issue by saying 'John is tall by academic standards but not basketball standards', and the dispute will disappear. By contrast, when two people disagree over 'Abortion is wrong', the dispute cannot be resolved in this way.

If correctness is understood as a sort of truth, then this behavior tends to suggest that the apparently clashing ordinary existence assertions by members of different communities should be understood as expressing different contents, with different truth-conditions. One might understand this difference as a sort of ambiguity in terms such as 'exist', or one might understand it as a sort of context-dependence associated with such terms.[4] On the first view, the speakers will be using two different words with different meaning, while on the second view, they will be using context-dependent terms with the same meaning (or character) in different contexts, yielding different contents. I think the second view is somewhat more plausible than the former, especially once one notes that the same phenomena can arise for different speakers and communities among the class of English speakers, and that postulating widespread ambiguity in 'exists' amongst this class is not especially desirable. But both views are available, as are views on which the truth-conditions of these assertions are invariant and only correctness-conditions are variable. I will return to the proper semantic treatment of ordinary existence assertions, outlining a potential contextualist analysis, later in the paper.

It is important to note that the analysis here is restricted to apparent disagreements that turn on differences in commonsense ontology. There are many disagreements involving ordinary existence assertions that go deeper than this. When two physicists disagree about whether there are six or twelve fundamental particles, or when two ornithologists disagree about whether the ivory-billed woodpecker still exists, their attitudes will be less relaxed than in the case above. But these are not cases that turn on differences in commonsense ontology, as I have characterized it, but rather on differences in first-order reasoning about the world. Later, I will give an analysis of ordinary existence assertions that account for the difference in character between apparent disagreements of each sort.

[4] At the same time, this behavior counts against a relativistic treatment of ordinary existence assertion, where these assertions would be taken to be assertions of a common content, which are true or false only relative to certain standards. Common contents are usually postulated by relativists in other cases precisely to account for the phenomenology of disagreement between speakers or communities. In the absence of this phenomenology, there is much less motivation for a relativist view, and a contextualist view is correspondingly stronger. See Egan, Hawthorne, and Weatherson (2005), Macfarlane (2005), and Stanley (2006) for more on relativism and contextualism.

4 Disagreements in Ontological Theory

We can now raise the same question for ontological existence assertions. Could standards of correctness for these assertions be variable in a similar way? I will take it that the correctness of an ontological existence assertion coincides with its truth. One can then raise the question by asking whether the truth-value of an existence sentence can vary between different ontological contexts of utterance, or between different ontological contexts of evaluation.

Consider an ontological disagreement between a nihilist and a universalist, faced with two particles in a vacuum chamber.

> Nihilist: There are exactly two objects in the chamber.
>
> Universalist: There are exactly three objects in the chamber.

Is this disagreement substantive or terminological? Some deflationists about ontology (e.g. Hirsch 2005) hold that the disagreement is terminological. For example, perhaps the two theorists mean different things by 'object': the nihilist means 'exists_n' (which applies only to simple objects), where the universalist means 'exists_u' (which applies to simple objects and to sums of simple objects). Or perhaps the two theorists mean different things by 'There is': the nihilist might mean 'There exists_n' (where 'There exists_n an X' is true iff there is a simple X), while the universalists means 'There exists_u' (where 'There exists_u an X' is true iff there are things arranged X-wise).

I am very much in favor of trying to resolve philosophical disputes into terminological disputes wherever possible. But in this case, I think the resolution does not succeed.[5] Here we can use a test for the presence of terminological disputes (see Chalmers forthcoming): does the dispute disappear once different senses for the problematic terms are distinguished, or does it persist as strongly as ever? In many cases, the dispute disappears, in which case we can diagnose the dispute as terminological. In other cases, the dispute is reduced but one finds some residual disagreement, in which case we might diagnose the original dispute as partially terminological and partially substantive. In still other cases, the dispute persists as strongly as ever, in which case the diagnosis of a terminological dispute is implausible.

In the case of ontological disputes, when one distinguises senses for the problematic terms, the disputes appear to persist as strongly as ever. For example, even once one distinguishes 'n-object'

[5] Bennett (this volume) gives another argument for the conclusion that disputes such as this one are not terminological.

and 'u-object', 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_n', 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?'

It may be that there is some ontological discourse that does not function in quite this way. Some "lightweight ontologists" might take a relaxed attitude toward apparent ontological disagreement. A sufficiently relaxed nihilist might hold that both the universalist and the nihilist are right relative to their own standards, and find nothing substantive left to argue over. But this would be an unusual attitude. A typical ontologist will find a substantive disagreement here, and will have no inclination to consider the dispute resolved once moves of this sort are made. Serious ontological dispute typically involves the sense that one is differing not just in classificatory practices but in substantive commitments. So I will take it that at least paradigmatic ontological existence assertions function in this way.

This phenomenon strongly suggests that whereas apparent ontological disagreement involving ordinary existence assertions is terminologically resolvable, apparent ontological disagreement involving ontological existence assertions is not. Perhaps there is something else that is wrong with these disputes, but the diagnosis of terminological disagreement does not seem to get to the heart of the matter. The structure of the discourse strongly suggests that participants do not use 'There exists' to express two different concepts, as happens in a terminological dispute. Rather, they use 'There exists' to express (or at least to attempt to express) a single common concept.

Furthermore, the phenomenon above tends to suggest that participants use 'There exists' to express a *primitive* concept. This is reflected in the fact that when we try to resolve the dispute, at each point the concept of existence (or a cognate concept) reoccurs. Here the dispute unlike a dispute over someone is a murderer, which might quickly resolve into disputes about whether they wielded the knife, and so on. It is more like a dispute (between a Kantian and a consequentialist, say) about what someone ought to do in a given situation. Here attempts to resolve the dispute will simply lead to reuse of the concept of *ought*, or cognate moral concepts such as *right*. Something similar applies to certain disputes over whether some organism is conscious. These expressions—'ought', 'conscious', 'exists'—appear to be especially resistant to analysis in more basic terms, at least on the relevant sort of usage. In effect, they function in these debates as primitive concepts.[6]

[6]For much more on this notion of primitive concepts and its relation to the resolution of disputes, see Chalmers

So there is at least prima facie reason to think that participants in ontological debates use terms such as 'exist' to express a single, common, primitive concept. We might call this the concept of *absolute existential quantification*. Since this concept, if it exists, is primitive, there may be little one can do to give an informative analysis of it in independent terms.[7] But one can at least say some things to elucidate the concept. Uses of the concept attempt to quantify over absolutely everything—that is, over everything that exists, in the most fundamental sense of 'exists'. Of course these elucidations are circular, but they may at least be helpful.

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). I will return to this matter later. For now I will talk about the concept of absolute quantification in order to elucidate the properties that it has if it exists.

5 Ontological Realism and Ontological Anti-Realism

We are now in a position to state ontological realism and anti-realism more precisely. Ontological realism, at least in its strongest variety, holds that every unproblematic ontological existence assertion has an objective and determinate truth-value. Or at least, it holds that every paradigmatic ontological existence assertion has an objective and determinate truth-value. I will set aside nonparadigmatic ontological existence assertions (of the sort discussed above) in what follows.

(forthcoming). It is worth noting that although it is plausible that all primitive concepts in this sense are unanalyzable, mere unanalyzability does not suffice for primitiveness in this sense. For example, there may be no good analysis of a term such as 'fish', but disputes over whether something is a fish do not behave in the manner characterized above.

[7] Absolute quantification should not simply be identified with unrestricted quantification. The absolute quantifier is plausibly intended to be an unrestricted quantifier, but it may be that some lightweight quantifiers of the sort used in ordinary existence assertions (see section 6) can also be regarded as unrestricted quantifiers, depending on one's views. It might be more viable to identify absolute quantification as unrestricted heavyweight quantification, at least if there is only one unrestricted heavyweight quantifier. Throughout this paper I pass over worries about set-theoretical paradoxes involving unrestricted quantification. For a discussion of these issues, see Rayo and Uzquiano (forthcoming).

We can say that an assertion has a determinate truth-value when it is has the truth-value "true" or "false". An ontological existence assertion has an objective truth-value if its truth-value does not depend on a context of utterance or a context of assessment: that is, if every ontological utterance of the same sentence has the same truth-value, and if the truth-value of these utterances do not vary with different ontological contexts of assessment.[8] An assertion is unproblematic when its nonquantificational vocabulary does not pose any obstacle to the assertion's having an objective and determinate truth-value: this requires (perhaps *inter alia*) that the nonquantificational vocabulary is truth-apt, not vague, not context-dependent, and not relativistic, and such that assignment of a truth-value to the sentence does not generate paradoxes akin to the liar paradox.

Ontological anti-realism is the denial of ontological realism. We can say that *weak* ontological anti-realism holds that not every unproblematic ontological existence assertion has an objective and determinate truth-value. We can say that *ultra-strong* ontological anti-realism holds that no unproblematic ontological existence assertion has an objective and determinate truth-value. I call the latter view "ultra-strong" because it is arguable that even a fairly strong anti-realism can allow that a limited class of ontological existence assertions have a determinate truth-value. In particular, there is much plausibility in the claim that ontological assertions of sentences such as 'There are odd perfect numbers', 'There are round squares', and 'There are unicorns' (or perhaps 'There are concrete unicorns', and so on), are determinately false, for first-order reasons largely independent of ontological reasoning. We might say that these assertions have a *trivial* truth-value, where the precise analysis of this notion will be left until later. One can then say that *strong* ontological anti-realism holds that no unproblematic ontological existence assertion has an objective, determinate, and nontrivial truth-value.

Of course there are many intermediate views. It is arguable that a fairly robust ontological realism is compatible with the claim that there is a limited class of unproblematic existence assertions with indeterminate truth-values. For example, consider a view according to which only particles and organisms exist (cf. van Inwagen 1990), while allowing that 'An X exists' has an

[8]There is room for disagreement about just what should be written into the definition of "objective" here. I am now inclined to think that context-independence shouldn't be written in: e.g. a Meinongian could think that 'exists' contextually expresses either being or existence, while still being an ontological realist. Instead views with context-dependent truth-values should be counted as versions of ontological pluralism, as discussed later, with a further distinction into heavyweight and lightweight ontological pluralism. Likewise, one could build mind-independence into the definition of "objective", but I choose not to on the grounds of the obscurity of that notion. On the current view, some views with mind-dependent truth-values will count as anti-realist, while others will count as lightweight realist, depending on the manner of the mind-dependence. (Thanks to Carrie Jenkins for discussion of these issues.)

indeterminate truth-value when X is a precise qualitative specification such that any entity satisfying that specification would be a borderline case of an organism.[9] One might reasonably consider such a view to be a version of ontological realism, although it deviates from strong ontological realism in the sense above. On the other end of the scale, we might consider a view on which that allows that fundamental particles exist determinately, and that there is no fact of the matter about the existence of nonfundamental entities. This view might be considered a fairly strong form of ontological anti-realism, although it deviates from strong ontological anti-realism in the sense above.

To capture intermediate views, one can also define ontological realism and anti-realism about specific domains. For example, ontological realism about numbers holds that all unproblematic existence assertions concerning numbers have an objective and determinate truth-value, while ontological anti-realism about numbers denies this. (Note that ontological anti-realism about numbers should be distinguished from the sort of anti-realism about numbers that denies that numbers exist—somewhat confusingly, this form of anti-realism about numbers is a form of ontological realism about numbers!) Likewise, one can define ontological realism and anti-realism about macroscopic objects, mereological sums, and so on. One might conceivably combine views, for example being an ontological realist about numbers but an ontological anti-realist about mereological sums.[10]

There are also different versions of ontological anti-realism depending on just how they think that the truth-values of unproblematic ontological existence assertions fail to be objective and determinate. An *ontological contextualist* holds that these truth-values are context-dependent. An *ontological relativist* holds that these truth-values are assessment-dependent. An *ontological indeterminist* holds that these truth-values are indeterminate. An *ontological noncognitivist* holds that ontological existence assertions are not truth-apt at all, perhaps because they function to express mental states other than beliefs. For reasons discussed earlier, I think that ontological contextualism is to be rejected (although a corresponding contextualism about ordinary existence assertions may be acceptable): if it were true, then many apparent ontological disagreements would be terminological, which I argued is implausible. Ontological noncognitivism also seems somewhat

[9] Van Inwagen himself holds that all unproblematic ontological existence assertions have a determinate truth-value. But for views that attempt to combine ontological realism with a limited degree of indeterminacy, see Koslicki (2004) and Korman (forthcoming).

[10] See Yablo (this volume) for a view that combines ontological realism about some domains with ontological anti-realism about others.

unmotivated. But this still leaves ontological indeterminism and ontological relativism on the table.

6 Lightweight and Heavyweight Realism

As I have stated it, ontological realism is compatible with views that are nevertheless quite deflationary about ontology. These include views on which ontological assertions have objective and determine truth-values, but on which the truth or falsity of these assertions is nevertheless shallow or lightweight. As before, we can call views of this sort versions of *lightweight ontological realism*, compared with *heavyweight ontological realism*, according to which the truth or falsity of ontological assertions is not lightweight in this way.

We can draw this distinction by first attending to an apparent feature of ordinary existence assertions. On the face of it, unconditional ordinary existence assertions such as 'There are prime numbers' are sometimes trivially correct, in that their correctness is knowable through trivial a priori reasoning. Likewise, ampliative existence assertions, such as 'If there are particles arranged heapwise, there is a heap' are often trivially correct. These assertions are ampliative roughly in that the consequent makes an existential claim that is not built into the antecedent. (That is, the consequent is not a logical consequent of the antecedent, where we take an expansive view of logical consequence such that for example, 'If x is a father, there exists someone who is an offspring of x' is a logical truth.)

One might well hold that some unconditional and ampliative ordinary existence assertions are *analytically* or *conceptually* true, if one holds that correctness is truth, and if one accepts the category of analyticity or of conceptual truth. It is not implausible, for example, that ordinary assertions of 'If there are particles arranged heapwise, there is a heap' are true in virtue of the way the terms are used, at least if any assertion is true in virtue of the way terms are used. If so, one could replace the appeal to triviality by an appeal to analyticity or conceptual truth. But even without these notions, it is plausible that these assertions are trivially correct in some related sense: for example, in that the thoughts they express can be justified by elementary a priori reasoning, yielding a priori knowledge. Here the thought expressed by an existence assertion is a thought that is true if the assertion is correct. I will leave the notion of elementary reasoning unanalyzed here, except to note that it is intended to exclude sophisticated philosophical reasoning.

We can say that an existence assertion involves *lightweight* quantification when it involves a use of existential language of a sort that can be used to express unconditional or ampliative exis-

tence assertions that are trivially correct. If one holds that existential language always expresses an existential concept, we can say that two uses are of the same sort when they express the same existential concept. If one holds that existential language does not always express an existential concept, one will have to characterize the relevant sorts of uses in some different way.

The discussion above suggests that some and probably all ordinary existence assertions involve lightweight quantification. From the fact that some unconditional or ampliative ordinary existence assertions can be trivially correct, it follows automatically that these assertions involve lightweight quantification. Furthermore, the uses of existential language in these assertions appear to be typical of the uses of existential language in ordinary existence assertions. For example, if these uses express concepts, the concepts used are typical of those concepts used in ordinary existence assertions. So there is good reason to suppose that all ordinary existence assertions involve lightweight quantification.[11]

By contrast, there is at least some plausibility in the claim that ontological existence assertions involve the use of *heavyweight* quantification: that is, they involve uses of existential language of a sort that cannot be used to express unconditional or ampliative existence assertions that are trivially correct.[12] Many ontologists hold that ontological assertions of sentences such as 'Numbers exist' are neither trivially true nor trivially false. Rather, their truth or falsity depends on substantive philosophical considerations that go well beyond elementary reasoning. The same goes for an ontological assertion of 'If there are particles arranged heapwise, there is a heap'. On the most common view, the consequent is not an analytic consequence of the antecedent, and is not otherwise a trivial consequence of it. Again, the truth or falsity of the consequent rests on substantive philosophical considerations that go well beyond elementary reasoning that starts from the antecedent.

On this view, positive unconditional ontological assertions are never trivially correct, and the only trivially correct conditional ontological assertions are non-ampliative. For example, the ontological assertions of the non-ampliative conditional 'If there exists an integer that is the sum of its proper divisors, there exists a perfect number' may be trivially correct, but ontological assertions of the unconditional 'There exists a perfect number' cannot be. Likewise, ontological assertions

[11] Some theorists will hold that lightweight quantification is not "real" quantification. Here I am using "quantification" in a broad sense to include language with the superficial appearance of quantification, but nothing turns on the terminological issue. One could equally talk in terms of lightweight and heavyweight quasi-quantification instead.

[12] For discussion of the question of whether ontologically relevant existence claims can be analytic, see Bennett (this volume), Dorr 2005, Field 1993, Hale and Wright 1992.

of the non-ampliative conditional 'If there exists an object with X and Y as parts and all of whose parts overlap X or Y, a mereological sum of X and Y exists' may be trivially correct, but ontological assertions of the ampliative conditional 'If X and Y exist, a mereological sum of X and Y exists' cannot be.

It is certainly not undisputable that ontological existence assertions involves heavyweight quantification, but the view is attractive. It is plausible that many ontologists at least take themselves to be using heavyweight quantification. If one sees ontology as an attempt to discover the fundamental structure of reality, then it is natural to hold that unconditional or ampliative ontological assertions are never trivially correct. And if one holds that the concept of absolute existential quantification is a primitive concept, as it at least superficially appears to be, then it is especially natural to hold that positive or ampliative claims involving this concept are never trivially correct.

In any case, we can now distinguish heavyweight and lightweight ontological realism more precisely. Heavyweight ontological realism, at a first approximation, is the conjunction of ontological realism with the claim that ontological existence assertions always involves heavyweight quantification. By contrast, lightweight ontological realism, at a first approximation, is the conjunction of ontological realism with the claim that ontological existence assertions always involve lightweight quantification.

As before, there may be intermediate views that still intuitively qualify as heavyweight or lightweight. For example, it plausibly suffices for a fairly robust heavyweight ontological realism to hold that *most* ontological existence assertions involve heavyweight quantification, and to hold the thesis of ontological realism for assertions of this sort (i.e., holding that all unproblematic ontological existence assertions of this sort have objective and determinate truth-values). Likewise, it might be compatible with lightweight ontological realism to hold that a small minority of ontological existence assertions involve heavyweight quantification, but that only those involving lightweight quantification have objective and determinate truth-values nontrivially.

Those ontological realists who hold that ontology is the study of the fundamental structure of reality are typically heavyweight realists. For example, Bennett (this volume), Dorr (2005), Fine (this volume), Horgan and Potrc (2006), Sider (this volume), and van Inwagen (this volume) are naturally read as heavyweight ontological realists. These theorists hold that unproblematic ontological assertions (usually or always) have an objective and determinate truth-value, and they hold that the truth or falsity of these assertions is a nontrivial matter to be determined by highly substantive philosophy.

By contrast, ontological realists who hold that ontology is largely a matter of conceptual anal-

ysis are typically lightweight realists. On this sort of view, the truth-value of ontological claims can be determined by conceptual reflection on the truth-conditions of existence claims, sometimes combined with first-order non-ontological knowledge of the world. Those who hold this sort of view typically hold that unconditional and/or ampliative ontological assertions can be analytic, in which case they are certainly trivial.

One verson of this view is the *commonsense realism* held by Hirsch (1993), who holds that ontology involves the analysis of ordinary concepts, and that our concept of existence in effect reflects the commitments of commonsense ontology. On this view, the truth of ontological existence assertions can be derived from qualitative truths and conceptual analysis. This view strays close to ontological anti-realism (and in particular to ontological contextualism), in that it allows that counterpart assertions (say, of 'There are cupcups') made by members of different communities with different commonsense ontological commitments may have different truth-values. But on this view, these assertions are not really using the notion of *existence* at all (which is tied to our ontology), but rather some different notion of *schmexistence*, and they should not be regarded as utterances of the same sentences as our existence assertions. So this view retains the letter of ontological realism, while sharing much of the spirit of ontological anti-realism.[13]

Other forms of lightweight realism include the *neo-Fregean realism* of Wright and Hale (2001, this volume), the *lightweight sortalism* of Thomasson (this volume), and the *lightweight maximalism* of Eklund (this volume). These closely related views are all liberal about the existence of objects, and hold that existence sentences can be conceptually analyzed in such a way that they can be analytically entailed by sentences without corresponding existence assertions. On the first view, ampliative conditions such as 'If the the Fs can be mapped one-to-one onto the Gs, there there is a number that is the number of Fs and the number of Gs' are taken to be conceptual truths. On the second view, all existential assertions have the underlying form 'There exists an F...' for some sortal concept F, and these sortal concepts have associated application-conditions such that the existential assertions can be analytic consequences of qualitative characterizations of the world. On the third view, it is taken to be a conceptual truth about existence that if the existence of an F is consistent with certain basic truths, then Fs exist. On all of these views, the truth-value of unproblematic ontological assertions is objective and (usually or always) determinate, so all can

[13] On a slightly different way of carving things up, one could lump together Hirsch's view and ontological contextualism as versions of (lightweight) *ontological pluralism*, and save "ontological realism" for views that do not admit multiple contents for 'exists' or its counterparts in assertions that function like ontological existence assertions in different communities.

be seen as versions of lightweight ontological realism.

It is also possible to be a sort of lightweight realist even if one rejects conceptual analysis and the notion of analyticity, and perhaps even if one rejects the notion of apriority. For example, a Quinean about analyticity and apriority might recognizably be a lightweight realist if he or she holds that ampliative existential conditionals (or inferences) are trivial, in the sense of not requiring substantive investigation. It is arguable that Quine himself was a sort of lightweight ontological realist, holding that ampliative conditionals of the form 'If Xs are endorsed by science, then Xs exist' are trivial.[14]

Lightweight realist views are all somewhat deflationary about ontology. In some ways, they have more in common with ontological anti-realism than with heavyweight ontological realism. In particular, they agree with strong ontological anti-realism that there are no nontrivial objective and determinate truths involving heavyweight quantification. The difference between lightweight realism and ontological anti-realism is in a certain sense semantic: the views differ mainly on their view of the content of ontological existence assertions. Both agree that *if* these are interpreted as involving heavyweight quantification, they do not have objective and determinate truth-values. Furthermore, the views can agree that if these assertions are interpreted as involving a certain sort of lightweight quantification, they have objective and determinate truth-values. The lightweight realist is automatically committed to this claim, and there is no obvious reason why the ontological anti-realist cannot allow that there is some possible lightweight quantificational expression that works this way.

This triangle between heavyweight realism, lightweight realism, and anti-realism is found in all sorts of areas of philosophy. One example is non-Humean realism about causation, a analytic regularity theory of causation, and an eliminativist theory of causation. Another example is property dualism about consciousness, analytic functionalism about consciousness, and eliminativism. In each case, the heavyweight realist gives inflationary truth-conditions and holds that they are satisfied, the lightweight realist gives deflationary truth-condutions and holds that they are satisfied, and the anti-realist gives inflationary truth-conditions and holds that they are not satisfied. The first and the second agree on the truth-value of certain sentences, while the first and the third agree on the truth-conditions of these sentences. The second and the third disagree on both of these linguistic matters, but consequently agree on the underlying character of the world: it is such that

[14]See Price (this volume) for more on the interpretation of Quine as a lightweight ontologist. I think that Frank Jackson and David Lewis are also most naturally read as lightweight realists, for whom 'If x and y exists, the sum of x and y exists' is a conceptual truth, although the reading of Lewis is not indisputable.

some parts of it satisfy the deflationary analysis, but no parts of it satisfy the inflationary analysis. In these cases, one can argue that the difference between lightweight realism and anti-realism is largely semantic.

Still, there are good and bad answers to semantic questions. In the following section, I will argue that typical ontological existence assertions involve heavyweight quantification, not lightweight quantification. If so, this favors ontological anti-realism over lightweight realsm.

7 In Favor of Ontological Anti-Realism

I accept ontological anti-realism, of a fairly strong kind. I think that ontological existence claims usually involve heavyweight quantification, but that claims of this sort usually lack a determinate truth-value, even if they are unproblematic in other respects. For example, I think that there is no fact of the matter about whether Platonism or nominalism is correct: ontological assertions of 'Numbers exist' lack a determinate truth-value. Likewise, there is no fact of the matter about whether universalism, nihilism, or some other view about the status of mereological sums is correct: ontological assertions of 'A mereological sum of X and Y exists' typically lack a determine truth-value.

My ontological anti-realism is qualified in one respect: I think that ontological existence assertions concerning certain fundamental entities (such as particles or people, perhaps) *may* be determinately true, although I am not sure of this. I will set aside this qualification for now, and return to it later.

Why ontological anti-realism and not lightweight realism? I think that it is plausible that *ordinary* existence assertions involve lightweight quantification, and I am inclined to accept a contextualist account of the content of the lightweight quantification that they involve. But we have seen that ontological existence assertions differ significantly from ordinary existence assertions, in that they involve the attempt to express a heavyweight quantifier: the absolute existential quantifier. I think one should take these attempts at face-value.

Lightweight realists often argue that the analysis of existence claims should be sensitive to the way they function in our discourse. They then give an analysis of the way these claims function in ordinary discourse, and use this to support lightweight realism. But once we make the distinction between ordinary and ontological existence assertions, it becomes obvious that ontological and ordinary claims function in our discourse in very different ways. So the arguments for lightweight analysis of existential claims may apply only to ordinary assertions and not to ontological asser-

tions.

In ontological discourse, reflective speakers typically have the strong sense that they are using an absolute quantifier, one that quantifies over what exists in the most primitive and fundamental sense, and one that is aimed at the fundamental structure of reality. There is at least the strong sense that there is a concept of absolute existential quantification, and that speakers are attempting to express this concept in their ontological existence assertions.

As far as I can tell, there is no bar to defining a quantificational expression, call it '$\exists_a$', such that this expression is stipulated to express the primitive concept of absolute existential quantification if such a concept is coherent (and such that the expression is defective if not). Reflective ontologists (see Dorr, Horgan, Sider, and others) tell us that this is the quantifier they intend to use. In ontological discourse, 'exists' functions very much in the way that one would expect '$\exists_a$' to function. So I think that there is good reason to accept that in ontological contexts, 'exists' has the content of '$\exists_a$'.

Furthermore, if there is an absolute quantifier that functions in this way, it is very plausible that it is a heavyweight quantifier. The absolute quantifier expresses a primitive concept, if it expresses any concept at all. Because of this, it is extremely implausible that ampliative conditionals involving the absolute quantifier, such as 'If x and y exist, the sum of x and y exists', or 'If there are particles arranged heapwise, there is a heap' could be analytic. It is unlikely that they are true in virtue of the concept of absolute quantification, because that concept is primitive and unanalyzable. It is unlikely that they are true in virtue of the concepts 'heap' and 'sum' alone, in part because they have logical consequences that do not involve these expressions. And it is unlikely that they are true in virtue of the concepts of absolute quantification and those expressed by 'heat' or 'sum' together: this combination might at best yield nonampliative analytic conditionals, such as 'If there is an object made of particles arranged heapwise, there is a heap', but not ampliative analytic conditionals. And if we move from analyticity to triviality, it is hard to see how an unconditional or ampliative conditional involving the absolute quantifier could ever be trivially true. So henceforth, I will take for granted that the absolute quantifier is a heavyweight quantifier.

As before, it may be that there are some existence assertions inside the "ontology room", made by sufficiently relaxed ontologists, that do not even attempt to use an absolute quantifier. Instead, these assertions may function much more like ordinary existence assertions, involving a lightweight quantifier. But as before, I am treating these assertions as nonparadigmatic, and restrict my claim to paradigmatic ontological existence assertions. I think it is highly plausible that at least some (probably most) ontological existence assertions attempt to express the absolute

quantifier, and it is these assertions with which I am concerned.

Why ontological anti-realism and not heavyweight realism? Both views can agree that paradigmatic ontological existence assertions express the absolute quantifier. The difference between our views is that I think the absolute quantifier is *defective*. Either it does not express a concept at all, or if it expresses a concept, that concept is defective too. In particular, the absolute quantifier does not have a determinate extension: something (a class of properties, say) that would combine with the extensions of otherwise unproblematic expressions to yield a determinate truth-value. Rather, if it has an extension at all, its extension is highly indeterminate.

I think it is plausible that there is at least a pseudo-concept of absolute quantification, one that functions in our ontological thought and discourse in many (though perhaps not all) of the same ways that a concept does. Whether there is a *concept* of absolute quantification depends on partly terminological issues turning on what is required of a "concept". For example, if concepts must have a reasonably determinate extension (or intension), then there may be only a pseudo-concept of absolute quantification and not a concept. If we are more liberal about what counts as a concept, then this will count as a concept, albeit a defective one.

The case of 'good' (and cognate moral expressions such as 'right' and 'ought') suggests that there mere fact that an expression functions as if it expresses a primitive concept does not entail that it has an objective and determinate extension. 'Good' functions in the same primitive way that 'exists' does. In fact, in the case of 'good', even ordinary and not just philosophical uses seem to function in this way. But this functioning is quite compatible with anti-realist views on which 'good' is defective, in that it does not yield a determinate truth-value for sentences containing it. Again, we could choose to say that on these views 'good' expresses a pseudo-concept, or that it expresses a defective concept. I will usually speak the latter way, both for 'good' and 'exists', but not much turns on this.

In the case of 'exists', the discontinuity between ontological and ordinary existence assertions gives us reason to think that the former involve absolute as opposed to lightweight quantification. But the same discontinuity should make us suspicious about whether we really have a nondefective grasp of the notion of absolute quantification. If ordinary practice involving 'exists' always involves lightweight quantification, then the coherence and nondefectiveness of this practive gives little support to the coherence and nondefectiveness of practice involving the absolute quantifier. It 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. If so, then one may reasonably have doubts about whether it has a determinate content.

In particular, one might question whether we really have a grip on what it would be for a table to "really exist" versus what it would be for a table to fail to exist, given that we are holding the lightweight existence of the table constant: that is, given that we are holding constant the underlying distribution of matter in virtue of an ordinary assertion of 'the table exists' is correct. Likewise, one might question whether we really have a grip on what it would be for numbers to exist, or to fail to exist, given that we are holding the lightweight correctness of mathematics constant. Our ordinary thought and talk of existence, as when we talk or think of a table's existing, or of prime numbers' existing, has conditions of correctness given by a lightweight quantifier, and there is little in our ordinary practice to suggest that we have a grip on truth-conditions of absolute quantification that transcend these lightweight conditions of correctness. And even on philosophical reflection, it is far from clear that we have any such grip.

Full-scale arguments against heavyweight realism must await another paper, but perhaps the simplest argument is the one at the start of this paper. Given full knowledge of the properties of two objects (including knowledge of the relations they bear to each other and to other objects, but not including any properties concerning relations to their sum), one is thereby in a position to trivially know *everything* about the two objects. Intuitively, there is no further nontrivial truth to resolve concerning whether the mereological sum of the two objects really exists. But if heavyweight realism were true, there would be such a further truth that would not be trivially knowable on the basis of this knowledge. So heavyweight realism is false.

Of course the main intuition in this argument is very close to a simple denial of heavyweight realism concerning mereological sums. But still, it is an intuition that many people (especially outside the field of ontology) will share. So this argument at least brings out a powerful intuition in favor of denying heavyweight realism.[15]

[15]Here are three more quick and closely related arguments, all turning on premises I have argued for elsewhere. *The creation argument*: In creating the world, God created the fundamental level. Once he did this, any further truths concerning the absolute existence of chairs or mereological sums were fixed automatically. But how were they fixed? Not by a contingent connection, as all truths supervene on the fundamental truths. Not by conceptual necessity, as the absolute quantifier does not enter into such necessities. And not by a pre-existing metaphysical necessity, as such a necessity must be a brute necessity, and there are no brute necessities. So there are no such further truths. *The conceivability argument.* If there is a nondefective absolute quantifier, then both nihilism and universalism are conceivable: neither of them can be reduced to contradiction by a priori reasoning. But what cannot be ruled out a priori is possible (setting aside Kripkean phenomena that are irrelevant here). Nihilism and universalism are not both possible. So there is no nondefective absolute quantifier. *The scrutability argument.* All truths are a priori entailed by fundamental truths (e.g. physical and phenomenal truths). But absolutely quantified sentences are not a priori entailed by fundamental truths. So there are no absolutely quantified truths. Note that strictly speaking, these arguments do not rule out the

8 Models, Worlds, and Domains

The central objection to ontological anti-realism, discussed earlier, goes as follows. Ontological anti-realism holds that ontological existence assertions can lack an objective and determinate truth-value, even when the nonquantificational vocabulary is not vague, context-dependent, or otherwise problematic. But the (absolute, unrestricted) existential quantifier is a logical constant, with a logically defined semantics that is objective and determinate. So this quantifier cannot fail to yield objectivity and determinacy for the assertion as a whole, when the other vocabulary is nonproblematic. So ontological anti-realism is false.

The response to this objection is as follows. It is true that logic gives a semantics for the existential quantifier. But this semantics tells us only how to evaluate a quantified statement in a *model*.[16] For assertions of quantified statements to have a truth-value, we have to evaluate them at a *world*. Worlds are not models. If so, the semantics of existential quantification does not yield a truth-value for quantified statements at worlds.

Models are a highly specific sort of abstract entity. They come with structure of various sorts, but what is most important for our purposes is that they come with a built-in *domain*. Intuitively, the domain is the domain of everything that exists with respect to that model. According to the standard semantics, an existentially quantified sentence $\exists x F x$ will then be true at a model if there exists an element of the domain such that the predicate F is true of that entity.

Worlds, on the other hand, do not obviously come with built-in domains. Of course the notion of a possible world is a technical notion within philosophy, and one might stipulate a notion on which worlds, like models, are abstract objects that come with built-in domains. But the real question is whether *this world*—the huge concrete reality within which we live—comes with a built-in domain, or at least with a canonically associated domain that serves as the domain of quantification for quantified statements. If it does not, then it does not determine a canonically associated model, and we cannot use the standard semantics to straightforwardly determine a truth-value for quantified sentences at this world.

We might say that the absolute quantifier requires an *absolute domain* for its evaluation. The ontological realist holds that the world has an associated absolute domain: the domain of every-

existence of absolutely quantified truths about the fundamental level. Of course there is an enormous amount more to say about each of these arguments.

[16]Here I use 'model' in the broad sense according to which all model structures are models, not just those model structures in which the sentences of a given theory are true.

thing that exists, in the most fundamental sense of 'exists'. By contrast, the ontological anti-realist denies that the world has an associated absolute domain.

Once things are cast this way, one can see that the ontological realist is committed to a very strong claim about the fundamental structure of reality. On this view, the fundamental structure of reality involves, or at least determines, an absolute domain of entities. By contrast, the ontological anti-realist holds that the fundamental structure of reality is less rich than this: it does not involve or determine an absolute domain of entities. The world may have structure of many sorts, but an absolute domain is not among that structure.

In any case, it is clear that there is no straightforward argument from logical semantics to the determinate truth of quantified statements. If anything, once things are viewed this way, it appears that the ontological realist is faced with a potential explanatory mismatch between semantic theory and metaphysics, one that requires a strong further commitment to resolve.

9 An Analysis of Ordinary Existence Assertions

The semantics of existential quantification still poses a challenge for ontological anti-realists, however. Assuming that ontological realists can make the case for an absolute domain, they can then give a straightforward treatment of ontological existence assertions: an ontological assertion of '$\exists x Fx$' is true if and only if an object in the domain has the property expressed by 'F'. Depending on their views, they may then be able to use this domain to play a role in giving correctness-conditions for ordinary existence assertion, perhaps via domain restrictions, or some other machinery.

By contrast, without an absolute domain, it is less obvious how an ontological anti-realist is in a position to analyze truth- or correctness-conditions for either ontological or ordinary existence assertions. In the case of ontological existence assertions, this might not be seen as too much of a cost, as the ontological anti-realists might deny that these assertions have truth-conditions or correctness-conditions at all. But there will remain the challenge of handling ordinary existence assertions, without the advantage of appealing to a built-in domain.

For ease of discussion, let us assume for now that correctness-conditions are truth-conditions, although not much will turn on this. The ontological anti-realist than has at least two options. One strategy would be to give nonstandard truth-conditions for ordinary existence assertions on which these do not involve quantification over a domain at all. Another strategy is to retain the tools of quantification over domains, by finding some way to associate every ordinary existence assertion

with a domain (perhaps a different domain for different assertions), which we can then use for semantic analysis. On the face of it, the second strategy is to be preferred if it is viable, as the semantic treatment of quantification in terms of domains is so powerful and familiar that much would be lost if we could not appeal to it.

Here, I will briefly spell out a version of the second strategy. In this section I will apply it to ordinary existence assertions. In the next section I will indicate how the strategy might be used to help analyze ontological existence assertions. In the section after that, I will address questions and objections. The proposal here should be considered highly speculative, and I am by no means certain that it is correct. It should not be read as offering an account of the logical form of existence sentences, and need not be read as offering an account of the propositions expressed by these sentences, or as a conceptual analysis of these sentences. Rather, it can be read as tool for helping to understand the conditions under which various sorts of existence assertions are true or correct, thereby shedding light on the commitments of ontological anti-realism as a whole.

The key idea can be seen as a relative of the familiar semantics whereby every quantified assertion is associated with a contextually-determined domain restriction function, picking out some subset of an overall domain as the domain of quantification. Instead of a domain restriction function, I will appeal to a contextually-determined *domain determination function*, which fixes the overall domain of quantification associated with an assertion.

To do this we can introduce some technical apparatus. Let us say that a *furnished world* is an ordered pair of a world and a domain.[17] Possible worlds can be understood in a variety of familiar ways. Intuitively, a domain is a catalog of entities that are taken to exist in a given world. To a first approximation, we might model a domain as a class of singular terms in an idealized language. More generally, we can model a domain as a class of ordered pairs $(\langle F_1, n_1\rangle, \langle F_2, n_2\rangle, ...)$, where each F_i is a predicate and n_i is a cardinal number. Intuitively, a domain will represent that there are exactly n_i objects satisfying F_i, for each i, and that every object satisfies F_i for exactly one i. It will also be required that the predicates F_i are sufficiently specific, in order that the world and the domain jointly fix the properties of the objects in question. I will later elaborate on how to make this intuitive constraint precise, with a notion of completeness for a domain, as well as elaborating a notion of equivalence between domains (so that distinct linguistic classes of this sort in effect specify the same catalog of entities, for example because they merely differ in order).

[17] Strictly speaking, I think this analysis should start with the notion of a furnished scenario, where scenarios correspond to a certain sort of maximal epistemic possibility rather than a maximal metaphysical possibility (see Chalmers 2004). This will yield an analysis of the primary intension of an ordinary existence assertion, whereas the furnished

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.

One may also impose a further constraint, holding that only certain furnished worlds are *admissible*. For a furnished world to be admissible, the domain must be an admissible domain for the corresponding world, where the notion of an admissible domain would need to be elaborated. For example, one might at least require that a admissible domain be internally consistent (no round squares), and that it be consistent with the qualitative character of a given world (no objects made of matter in a world without matter). We could then say that an admissible furnishing function is one that maps a world to a domain that is admissible for that world, so that a world and an admissible furnishing function yield an admissible furnished world. This constraint will not be crucial in this section, but I will elaborate on it later.

Intuitively, a furnishing function is a technical counterpart of a Carnapian ontological framework. Different ontological views will correspond to different furnishing functions. There will be a nilihist furnishing function, mapping worlds to domains representing a class of simple entities that are taken to exist on a nihilist view. There will be a universalist furnishing function, mapping worlds to domains representing a class of simple and complex entities that would be taken to exist on universalist views. There will be furnishing functions that admit abstract objects, and furnishing functions that admit only concrete objects. There may be a commonsense furnishing function that admits those entities taken to exist according to commonsense ontology. And so on.

With furnishing functions in hand, we can then postulate that for any ordinary utterance, the context of utterance determines a furnishing function. Intuitively, this function corresponds to the ontological framework endorsed by the speaker in making the utterance. For example, ordinary discourse about tables and chairs may involve a context that determines a commonsense furnishing function. Typical mathematical discourse may involve a context that determines a furnishing function that admits all sorts of abstract objects, and so on.

If we make the standard assumption for every utterance there is a world of utterance, then the world of utterance combined with the furnishing function of the context of utterance will together determine a domain. It is this domain that will be used to assess the truth of the utterance.

Here we can appeal to a variant of the familiar idea most linguistic expressions (or utterances

world analysis yields a secondary intension. But to keep the discussion reasonably simple, I will ignore this complication.

thereof) determine functions from possible worlds to extensions. For example, a singular term might determine a function from worlds to individuals in those worlds, while general terms and predicates might determine a function from worlds to classes of entities in those worlds, and so on. We can put forward the following variant of this idea: most linguistic expressions (or utterances thereof) determine functions from (admissible) *furnished* worlds to extensions. A singular term will determine a function from furnished worlds to individuals (here represented as entities in the domain of the world), general terms and predicates will determine functions from furnished worlds to classes, and so on.

On the face of it, this view is reasonable. One reason to think that expressions determine functions from possible worlds to extensions is that our linguistic competence allows us to evaluate expressions under various suppositions about the world: if the world is like so, then Jack the Ripper is such-and-such. (There is also a counterfactual version: if the water were like so, then water would be such-and-such. Both versions can be accommodated under a two-dimensional semantics, but I will not make much of the difference here.) But our competence also allows us to evaluate expressions under different suppositions about which entities exist in the world, even holding the qualitative nature of the word constant. For example, we can say that if nihilism is true, then Jack the Ripper does not exist, and so on. So it is not a great stretch to allow that our expressions determine functions from furnished worlds to extensions.

If this is so, then applying standard compositional semantics, non-quantified sentences (or utterances) will determine a function from furnished worlds to truth-values. We can then use the standard semantics for evaluating an existentially quantified sentence (or utterance) at a furnished world: it is true if the corresponding open sentence is true of some entity in the domain of that world. By the usual semantics, an open sentence is akin to a predicate, determining a function from worlds to classes. So by the semantics above, an open sentence will determine a function from furnished worlds to classes. An existentially quantified utterance will be true at a furnished world w if the function associated with the corresponding open sentence, evaluated at w, yields a nonempty class. Likewise, a universally quantified utterance will be true at a furnished world if the value of the function at the world is the whole domain, and so on. Here I have ignored the effects of ordinary quantifier domain restriction, but one can integrate these into the picture in a straightforward way.

We can then say: an ordinary utterance is correct at a world w iff it is true at the furnished world $\langle w, f(w)\rangle$, where f is the furnishing function specified by the context of utterance. An ordinary utterance is correct (*simpliciter*) iff it is true at the world of utterance. If the correctness

of an ordinary utterance coincides with its truth, as I think it does, then we these definitions will also give conditions for the truth of an ordinary utterance at a world, and its truth simpliciter.[18]

In this way, we can see that ordinary utterances determine ordinary truth-conditions across possible worlds. Of course an utterance of the same sentence by speakers whose context of utterance determines different furnishing functions will have different truth-conditions. For example, an utterance of 'There are cupcups' may be true in a context with a liberal furnishing function, and false in a context with a less liberal furnishing function. Likewise, in the case discussed earlier in the paper, an utterance of 'There are two objects on the table' by a human may be true, when an utterance by a Martian in an otherwise identical context may be false, simply because of a difference in their furnishing functions.

This account is neutral on how easy it is for furnishing functions to vary between contexts, and how widespread such variations are. In principle, the account is compatible with the view that all utterances within a community are associated with the same furnishing function, where this corresponds intuitively to the ontological framework endorsed by the community. So, for example, utterances in our community might be associated with a "commonsense ontology" furnishing function, while utterances in a Martian community might be associated with a universalist furnishing function. But once contextual variation is admitted, it is natural to hold that variation is more widespread than this, so that utterances by speakers in the same community might be associated with different furnishing functions, and likewise for utterances by the same speaker in different contexts.

Of course there may be considerable vagueness as to exactly which furnishing function is determined by the context of utterance, just as there may be vagueness as to exactly which domain restriction is determined by the context of utterance. This vagueness will produce vagueness concerning the truth-value of the utterance as a whole. The vagueness might be handled by one of the many standards tools for dealing with vagueness: for example, by supervaluating over numerous furnishing functions, by introducing fuzzy or vague furnishing functions, by postulating an unknown fact about the furnishing function, and so on. The problems here do not seem to go

[18] A better-elaborated version of this picture will allow for the possibility that more than one furnishing function can be associated with an utterance. Here one can appeal to an analog of the analysis of quantifier domain restriction according to which every determiner phrase is attached to a domain restriction. A natural idea is that every determiner phrase is attached to a furnishing function. More generally, it is natural to hope that a number of ideas that have been developed in the analysis of domain restriction can be extended to counterparts in the analysis of domain determination, although the viability of such extensions requires significant further work.

beyond the sort of problems already introduced by vagueness of quantifier domain restrictions, and other sorts of vagueness.

10 An Analysis of Ontological Existence Assertions

What about ontological existence assertions? Can we give any sort of truth-conditional analysis of these? The ontological anti-realist could hold that ontological existence assertions are so hopelessly defective that they do not express propositions or have truth-conditions at all. But there is room for a milder view. For example, we have seen that an ontological anti-realist might allow at least that ontological assertions of 'There are concrete unicorns' and the like can be false. And there might be views that are ontologically anti-realist about certain domains, while being realist about others. Here it would be useful to have some sort of truth-conditional analysis of what is going on on these views.

One view might use the same semantics as the ontological realist: ontological existence assertions are true iff the corresponding open sentence is true of some entity in the absolute domain. This view could then be combined with the thesis that in our world, the absolute domain is massively indeterminate. It may be determinate that certain things are not in it—that is, that certain expressions do not represent entities in it. It might even be determinate that certain entities are in it. But for a very wide class of expressions, it is indeterminate whether these have a referent in the absolute domain or not. This view is not out of the question, although it does require that some sense can be made of the notion of an absolute domain, which an ontological anti-realist might question.

However, this view suggests an alternative treatment, one that is compatible with the view just mentioned but does not require it. As already mentioned, it is common to understand indeterminacy using *supervaluation*. For example, on some views of vagueness, a use of an vague predicate such as 'tall' is associated with a range of properties that are admissible extensions, such that 'John is tall' is true iff John has ϕ for all admissible extensions ϕ (of this utterance of 'tall', false is John has ϕ for no admissible extension, and indeterminate otherwise.

We can apply this idea to ontological existence assertions as follows. We have already seen that some furnishing functions are admissible and some are not. We can then suggest: an ontological existence assertion $\exists x Fx$ is true at a world w iff it is true at the furnished world $\langle w, f(w)\rangle$ for all admissible furnishing functions f. The assertion is false at w iff it is false at $\langle w, f(w)\rangle$ for all admissible f. Otherwise, its truth-value is indeterminate at w.

Of course this requires us to say more about what it is for a furnishing function to be admissible, which requires us to say more about what it is for a domain to be admissible for a world. I will say more about this in the next section. But intuitively, there should at least be a consistency constraint: for example, if w contains no matter at location l, then the domain cannot contain an entity specified as being made of matter at l. There *might* also, depending on the strength of one's ontological anti-realism, be positive constraints concerning certain entities that must be in any domain that is admissible at a given world. For example, a medium-strength ontological anti-realist might hold that at our world, fundamental particles and/or people must be part of any admissible domain. There could also be negative constraints concerning entities that cannot be in any admissible domain.[19]

As long as more than one domain is admissible at a world, then some unproblematic ontological existence assertions will have indeterminate truth-values. On the strong anti-realist view, the only constraint on admissibility will be the consistency constraint, so that (for example) nihilist domains, universalist domains, nominalist domains, and Platonist domains will all be admissible, as will the corresponding furnishing functions. In our world w, 'Tables exist' will be true at $\langle w, f(w)\rangle$ for some furnishing functions, and false at $\langle w, f(w)\rangle$ for others. So an ontological assertion of 'Tables exist' will be indeterminate at our world. The same goes for 'Numbers exist', and so on.

Which statements will have determinate truth-values? On an ultra-strong anti-realist view, there may be no constraints on admissibility (not even consistency constraints), so no ontological assertions will have determinate truth-values. On a strong anti-realist view, there will be consistency constraints, so that that domains with round squares may be inadmissible at any world (because of internal inconsistency). If so, 'There are round squares' will be false at $\langle w, f(w)\rangle$ for any world w and any admissible furnishing function f, so an ontological assertion of the sentence will be false at all worlds. Likewise, domains with unicorns may be inadmissible at our world @ (because of inconsistency with the specification of our world), so 'There are no unicorns' will be false at $\langle @, f(@)\rangle$ for all admissible f, so an ontological assertion of the sentence of 'There are no unicorns' will be true at our world.

On intermediate anti-realist views, there may be other positive and/or negative constraints on admissibility, so that for example all admissible domains for a world must contain certain

[19] As long it is antecedently plausible that it is not vague whether a furnishing function is admissible, this framework will not have a problem of higher-order indeterminacy analogous to the problem of higher-order vagueness in supervaluationist treatments of vagueness.

fundamental entities in a world, or so that no admissible domains can contain abstract objects. If so, it will may be that ontological assertions of 'There are fundamental particles' will be true at our world, and/or that ontological assertions of 'There are numbers' will be false at our world, while other ontological assertions will have an indeterminate truth-value.

On a strong ontological realist view, on the other hand, there will be exactly one domain that is admissible for every world (at least up to equivalence between domains): the absolute domain for that world. There will likewise be just one admissible furnishing function, mapping every world to its absolute domain. On this view, an ontological assertion will be true at a world w if it is true at $\langle w, f(w)\rangle$, where $f(w)$ is the absolute domain of w. This will then yield the same results as the semantics of existence assertions, where in effect worlds come with built-in domains. On this view, every unproblematic ontological assertion will have a determinate (and objective) truth-value.

11 Questions and Objections

Of course this framework immediately raises many questions. I will address some of them in what follows, in a somewhat random order. Readers should feel free to pick and choose.

What is admissibility?[20]

An immediate question is to say more about what it is for a domain to be admissible at a world. I will focus first on the consistency constraint. Here the idea was that the specified domain must at least be consistent with the world. Intuitively, given a world with two cups, domains both with and without a cupcup are consistent with the world. But in a world with no ectoplasm, domains containing entities made of ectoplasm are inconsistent with the world.

To make the notion precise, we need to first be more precise about domains. As before, a domain is a class of ordered pairs $(\langle F_1, n_1\rangle, \langle F_2, n_2\rangle, ...)$, where each F_i is an unproblematic predicate and n_i is a cardinal number. We can say that an *existence sentence* for an ordered pair $\langle F_i, n_i\rangle$ is a sentence saying that there exist exactly n_i objects satisfying F_i. A *totality sentence* for a domain is a sentence saying that every object satisfies F_i for some i and that no object satisfies F_i and F_j where $i \neq j$. A *domain sentence* for a domain is a conjunction of existence sentences for each ordered pair in the domain with a totality sentence for the domain. Here we assume an

[20]This section is still a mess. Feel free to skip it if you're reading before morning coffee.

idealized infinitary language with unlimited expressive power throughout.

A domain is then admissible at a world if its domain sentence is consistent with that world.[21] We can say that a domain sentence is consistent with a world iff it is consistent with a *world-sentence* for that world (or with all world-sentences for that world, if there is more than one). A world-sentence for a world w is an unproblematic sentence that is true at w and at no other worlds. One might additionally require that a world-sentence is a specification of the *fundamental* truths at w: for example, if physicalism is true at w, these might be microphysical truths and a that's-all clause. I will discuss this further in the next sentence. For now, I will be neutral on the character of a world-sentence, and I will also be neutral on whether or not it can include quantificational vocabulary.

It remains to define the consistency of two sentences. We cannot simply say that S_1 and S_2 are consistent iff $S_1 \& S_2$ is metaphysically possible. If S is a world-sentence, there is no sentence T such that $S \& T$ and $S \& \neg T$ are metaphysically possible (if there were, S would be true in more than one world). But we want to say that if S is a world-sentence for this world, S is consistent with both 'There is a cupcup' and its negation. So consistency is somewhat weaker than metaphysical compatibility.

We can instead suggest that S_1 and S_2 are consistent when they are logically compatible with a certain subset of necessary truths. To say that sentences are logically compatible is to say that no contradiction can be derived from them in a standard logical framework, in which logical rules for the quantifiers are exhausted by the standard introduction and elimination rules and interdefinitions. Intuitively, the relevant necessary truths are *ontologically neutral* necessary truths—those that make no commitments about ontology. In particular, any unconditional or ampliative existence sentences are excluded (e.g. 'Numbers exist', 'If X and Y exists, their sum exists'), as

[21] We should also require that a domain is complete at a world. A domain d is complete at a world w if for any two domains d_1 and d_2 such that d_1 and d_2 are both refinements of d that are consistent at w, d_1 and d_2 are consistent with each other. A domain $(\langle G_1, m_1\rangle, \langle G_2, m_2\rangle, \ldots)$ is a refinement of a simple domain $(\langle F, n\rangle)$ iff $m_1 + m_2 + \ldots = n$ and each G_i is a conjunctive predicate one whose conjuncts is F_i. Then a domain D is a refinement of a complex domain $\langle F_1, n_1\rangle, \langle F_2, n_2\rangle, \ldots\rangle$ iff it can be partioned into domains D_i such that for all i, D_i is a refinement of $\langle F_i, n_i\rangle$. Two domains are consistent with each other when their domain sentences are logically consistent.

We can also use these notions to define the equivalence of two domains at a world. Two domains D_1 and D_2 are equivalent at w when both are consistent at w and all refinements of D_1 are consistent with all refinements of D_2. In effect, equivalent complete domains specify the same catalog of entities at a world. In a fuller treatment, it is probably better to reidentify domains with equivalence classes of complete domains in the previous sense, and to regard furnished worlds as ordered pairs of worlds with domains understood this way.

are analogous sentences involving singular terms that logically entail such sentences (e.g. '7 is prime'). Also excluded are indeterminate existence claims, such as 'It is indeterminate whether numbers exist', and 'If X and Y exist, it is indeterminate whether their sum exists'. (On an ontological anti-realist view, many claims of the latter sort will be necessary.) Allowable truths will include nonampliative conditional necessary truths, or those whose necessity derives from non-quantificational terms: for example, 'If someone is a bachelor, they are unmarried' (which derives from the predicate 'bachelor'), perhaps 'If necessarily P, then necessarily necessarily P' (which derives from 'necessarily'), and even 'If Joe is a father, there is someone who is Joe's offspring' (which derives from 'father').[22]

A rigorous characterization of ontologically neutral necessary truths remains an open question (suggestions gratefully received!). One tricky issue is a precise definition of "ampliative" and/or the associated notion of necessity deriving from a nonquantificational term. A related issue concerns whether and how to exclude negative existence claims, such as 'Numbers do not exist'. Here the challenge is to find a way to exclude these, while allowing necessities such as 'There are no round squares'. Intuitively, the former derives (if true) from substantive ontology, while the latter does not, but it is not straightforward to make the distinction precise. An ontological realist can circumvent this issue on the grounds that 'Numbers do not exist' (and the like) are indeterminate rather than necessary on their view. But the cost is then that indeterminacy of 'Numbers do not exist' will then be used to explain why domains with and without numbers are admissible, so that the latter fact about admissibility cannot be used to explain the former fact about indeterminacy (as one might have hoped that the semantic analysis in the previous section might do).

In any case, assuming that the notion can be made precise, we can then use it to analyze consistency. For example, 'There are two cups' is consistent with both 'There is a cupcup' and its negation, as there are no ontologically neutral sentences with which the relevant pairs are inconsistent. So one can expect that a world-sentence for our world will be consistent with both 'There is a cupcup' and its negation. Correspondingly there is reason to expect that domain sentences entailing 'There is a cupcup' and domain sentences entailing 'There are no cupcups' will both be consistent with this world-sentence. So domains both with and without cupcups will be admissible at our world.

On the other hand, 'There is no ectoplasm' is not consistent with 'There is an object made of

[22]One can likewise have a version of this framework where S_1 and S_2 are consistent when they are logically compatible with a certain subset of ontologically neutral *a priori* truths. This notion can then be applied to the analysis of furnished scenarios, and in particular to determine which domains are admissible relative to a scenario.

ectoplasm', as these are inconsistent with the existence-neutral necessary truth 'If there is an object made of ectoplasm, there is ectoplasm'. So given a world whose world-sentence entails 'There is no ectoplasm', a domain involving objects made of ectoplasm will be inadmissible. Likewise, 'There are round squares' is inconsistent with the existence-neutral necessity 'If something is round, it is not square'. It is consequently inconsistent with any world sentence. It follows that domains involving round squares will be inadmissible at any world.

Finally, it will be observed that *if* the world-sentence of a world includes or entails any existence claims, such as 'There are photons', then sentences such as 'There are no photons' are automatically inconsistent with the world-sentence. If so, then domains without photons will be inadmissible, so that an ontological assertion of 'There are photons' will be determinately true at the world in question. I will look into this matter more in what follows.

What are the fundamental truths?

In the previous section, I appealed to the idea that worlds can be described by world-sentences. Going along with this is the idea that various existential claims are true, false, or indeterminate in virtue of their relations to such a world-sentence. So the world-sentence is presumably to be something more basic than those existential claims. One can naturally take it to be a characterization of the fundamental truths about a world. But what are these fundamental truths? And must they themselves involve existential quantification?

An ontological anti-realist can take various attitudes here. A hardline deflationist about metaphysics might reject the claim that some characterizations of the world are more fundamental than any other. Various sentences are true at only this world, and that is all there is to say. This would still leave the question of what sort of sentences *are* determinately true at a world, which would raise some of the issues in what follows, but it would avoid some of the others.

Still, nothing in the motivation for ontological realism so far described requires this hardline deflationism, and I am inclined to reject it. Intuitively, some contingent truths about the world hold *in virtue* of other contingent truths about the world. Truths about chemistry may hold in virtue of truths in physics, for example. We can then say that a *fundamental* truth about the world is a contingent truth that does not hold in virtue of other truths. It does not automatically follow from the earlier claim that there are fundamental truths—there could be an ever-descending sequence of truths, each of which holds in virtue of more fundamental truths (see e.g. Schaffer 2003). But under certain natural assumptions, there will be an end to such sequences, even if these are only

truths about the world as a whole. If so, it is then plausible that conjoining sufficiently many fundamental truths will yield a world-sentence for our world: one that is true only at our world.

One can then ask, what is the character of fundamental truths in our world, and in other worlds? At this point, a realist may naturally suggest that these truths are truths about objects and their properties: for example, about fundamental particles, or people. Such truths will involve either an existential quantifier, or singular terms, and either way they will entail some existentially quantified claims. Does this not lead to a degree of ontological realism?

Here the anti-realist has various options. They may claim that fundamental truths are not truths about objects at all. For example, they may hold that fundamental truths are truths about the distribution of stuff, with quantifiers for stuff rather than objects: there is some matter distributed with such-and-such densities at such-and-such locations, or there is some experience distributed in such-and-such way. This view gives up on ontological realism about objects, but it retains a variant of ontological realism concerning stuff. That is, in effect it will hold that counterparts of ontological existence assertions involving *stuff* quantifiers are determinately true, even if those involving object quantifiers are not. Of course if the reasons for denying ontological realism concerning objects also apply to ontological realism concerning stuff, then this view will be unsatisfactory. But it is not obvious that these reasons extend to stuff in this way, so this view is at least a contender.

Alternatively, an anti-realist may hold that fundamental truths are specified in an object-free predicate functor language where every sentence is analogous to 'It is raining', or 'Raineth'. For example, they may include 'Particleth' where someone else might have had 'There is a particle'. Or less artificially, a specification of the quantum wave-function of the universe might be regarded as making a claim of this sort. This view is formally clean, but it arguably leaves the character of the fundamental level of reality more obscure than most of the alternatives.

Another option is to accept that fundamental truths involve object-property language, but only lightweight object-property language. For example, they may include 'There is a photon at such-and-such spacetime location', 'There is a person with such-and-such experiences', where the existential quantifier is interpreted as a lightweight quantifier, such as the maximalist lightweight quantifier. Or perhaps one could make a lightweight claim 'The universe has such-and-such properties'. One cost here is that one must then give up the claim that lightweight quantified truths always hold in virtue of underlying truths. Another is that if we analyze admissibility in terms of consistency with fundamental truths, where these truths involve lightweight quantification, then one cannot use the notions of admissibility or consistency in analyzing the corresponding lightweight quantifiers. So this view may lead to a view on which one or more lightweight quantifiers (the maximalist or

nihilist quantifier, perhaps) is taken as conceptually primitive.

Perhaps the view most in the spirit of ontological anti-realism holds that there is more than one equally good fundamental description of the world, and that (for example) fundamental descriptions of any of the three sorts given above are allowable. On this view, reality has a fundamental nature, but this fundamental nature can equally well be described in terms of objects and properties (using a lightweight quantifier), in terms of stuff, or in terms of predicate functors. On this view, one can use any one of these specifications to state truth-conditions for sentences of the other sorts: using stuff language to state truth-conditions for lightweight quantified sentences, perhaps, or vice versa. But there will not be one such specification that is metaphysically fundamental, and one will not be able to reductively analyze all such sentences simultaneously. Of course this is not to say that anything goes in a fundamental description: for example, a physicalist will require that all of these descriptions must be in a broadly microphysical language, rather than a mental language, for example.[23] So on this view, reality will have a determinate fundamental nature, but a nature that is less fine-grained than some metaphysicians think.

A much more moderate anti-realist might hold that fundamental truths involve *heavyweight* object-property language. For example, these truths may include 'There is a photon at such-and-such location' or 'There is a person with such-and-such experiences' where the existential quantifier is interpreted as a heavyweight quantifier. Or perhaps one could make a heavyweight claim 'The universe has such-and-such properties'. On this view one must accept that heavyweight quantification is coherent, and that we possess a concept of absolute quantification with a nontrivial but not fully determinate extension. In addition, to avoid circularity, such an anti-realist will have to deny that the truth-conditions of heavyweight quantification are grounded in supervaluation over admissible furnishing functions. So one will have to find some other explanation of the indeterminacy of ontological existence assertions, though one will still be able to use the framework of furnishing functions as a useful tool.

On such a view, one gives up strong ontological anti-realism for an intermediate anti-realism that admits the heavyweight existence of certain entities. But importantly, this claim will be restricted to fundamental entities. We saw earlier that the considerations in favor of ontological anti-realism do not apply nearly as clearly to fundamental entities as to nonfundamental entities. And I think there is some attraction to the idea that we have a grip on the idea of a fundamental

[23]Carnap (1950) seems to adopt an even stronger anti-realism here, holding that there is no fact of the matter as to whether a physicalist or a phenomenalist view of the fundamental level is correct. Here my own anti-realism reaches its limits: I find it impossible to believe that *this* is something about which there could be no fact of the matter.

entity existing absolutely, even though we do not have the same grip on the idea of a number, a table, or a mereological sum existing absolutely.

Which view is best? I am not sure. I am perhaps most inclined to accept the view involving multiple fundamental specifications of reality, of which some involve lightweight quantification but none involve heavyweight quantification. But I do not yet exclude a view on which there are fundamental specifications of reality that involve heavyweight quantification (perhaps alongside other fundamental specifications that do not). The ontological anti-realism that I have advocated is perhaps most naturally developed into a strong anti-realism, but a moderate anti-realism that involves realism about fundamental entities is still on the table. Either way, I take the question of which specifications of reality are fundamental specifications to be among the hardest and most important questions in philosophy.

In any case, the results of this analysis cohere with the results of a related discussion by Sider (this volume): there is no detour around the whole of metaphysics. Even if one dismisses claims about the existence of numbers or mereological sums as indeterminate, one still needs to face up to hard metaphysical questions. The question of ontological realism should be distinguished from the more general question of metaphysical realism: where the former focuses on the existence of objects, the latter focuses on the nature of reality. Nothing in the considerations I have raised recommends a global skepticism about metaphysics. If anything, it suggests that metaphysical inquiry should be especially focused on questions at the fundamental level.

Is the use of abstract objects legitimate?

In this paper I have appealed frequently to all sorts of abstract objects, such as possible worlds, ordered pairs, functions, sentences, and so on. It might thought that this commits me to their existence, and therefore to rejecting ontological anti-realism at least in these domains.

In response: ontological realism about these entities is not required. When I have discussed these entities, I have been making ordinary existence assertions, not ontological existence assertions. When I have done so, I have been working within a liberal framework with a furnishing function that admits all sorts of abstract objects. In fact, whenever I do philosophy (and especially the philosophy of language), I work within such a framework.

Just as it is hard to do mathematics without appealing to numbers, it is hard to do philosophy without appealing to abstract entities. If one did not work within a liberal framework, many of the ideas in this paper would be much harder to state, and perhaps some of them would be impossible

to state, although this is not obvious. Likewise, if one did not work within a liberal framework that admits numbers and other mathematical entities, many mathematical results would become harder and perhaps impossible to state, as perhaps would many scientific theories. However, nothing here commits one to ontological realism.

Perhaps there is some indispensability argument that starts from the premise that appeal to abstract objects is indispensable in various areas of science, mathematics, and philosophy, and concludes that ontological realism is true of abstract objects. At the moment, however, I cannot see how such an argument would go. Such an argument might reveal the indispensability of working within a liberal framework in order to do science, mathematics, and philosophy. But that is a very different conclusion.

In fact, when doing philosophy it is often sensible to assume a maximalist framework, on which any entity whose existence is consistent with the nature of this world can be taken to exist (see Eklund, this volume). This makes for considerable convenience, not least in the present project. For example, from within this framework, one can drop the construction of domains as classes of linguistic entities, and simply regard them as classes of entities (simpliciter) in worlds. Furthermore, one can then regard furnishing functions as a sort of domain restriction function: they map a world to a subclass of the entities that exist in that world. This allows a more straightforward semantic treatment. The domain restriction function that performs the work of a furnishing function will still play a role that is conceptually separable from that of more familiar domain restriction functions, but one can use many of the same tools of analysis.

Again, nothing here entails maximalism as a heavyweight ontological view. It simply reflects the advantage of (lightweight) maximalism as a framework for conducting one's theorizing about the world.

Is this ontological pluralism?

Hirsch (2002), Eklund (this volume) and Sider (this volume) characterizean various deflationary views as involving *quantifier variance* or *ontological pluralism*: roughly, the view that there are many candidate meanings for the existential quantifier (or for quantifiers that behave like the existential quantifier in different communities), with none of them being objectively preferred to the other. Is this a view of this sort?

Yes and no. It is true that on this view, there are many ways to use quantifiers, to express claims with very different truth-conditions. In particular, ordinary uses of quantifiers can express quite

different contents. On the other hand, there is a distinguished use of the quantifier, in ontological uses, to express (or to attempt to express) the concept of absolute quantification. Insofar as this use has a content, it is a sort of privileged content.

On the account I have presented, there are not many different quantificational expressions that correspond to these contents. There is just one expression 'exists', and various cognate expressions, all of which have different contents in different contexts. However, there is no obstacle to defining new expressions that express each of these contents in a context-independent way. For example, I have already occasionally used the expression 'absolutely exists' as a context-independent expression for absolute quantification. One could also define terms such as 'exists_n', 'exists_u', 'exists_c', 'exists_m', and so on, to express the content that 'exists' has when used in contexts involving nihilist, universalist, commonsense, and maximalist frameworks, for example. So we would then have a plurality of quantificational expressions, although again there would be a privileged expression corresponding to the ontological use.

One might wonder about how these expressions would interact with contexts of use: for example, what happens if one uses 'exists_u' in a nihilist context? The answer is that as a non-context-dependent expression, 'exists_u' will have the same content that it has in a universalist context. In any context, '$\exists_u x Fx$' will be true at a world w if '$\exists x Fx$' is true at the furnished world $\langle w, f(w)\rangle$, where f is the universalist furnishing function. This expression is equally available in a nihilist context, although of course in such a context 'exists_u' will have a different content from 'exists'.

What are the substantive disagreements?

Once one explicitly distinguishes all the quantifiers above, it is possible to prescind from semantic issues about the English term 'exists' and examine the residual substantive agreements and disagreements between ontological anti-realism and lightweight and heavyweight ontological realism more directly. Most ontological anti-realists and lightweight realists can agree that all the lightweight quantifiers 'exists_u', 'exists_n', and so on can be defined. Some ontological realists can also agree to this, although not all will.[24] Those who agree that the lightweight quantifiers can be defined will agree about the truth-values of statements with these quantifiers, and about their epistemic status. For example, if M is 'is a mereological sum of distinct entities', then all agree that '$\exists_u x Mx$' is true in our world but '$\exists_n x Mx$' is false. Furthermore, if $S(a, b)$ stands for the sum of a and b, all will agree that '$\exists_u S(a, b)$' follows trivially from '$\exists_n a \& \exists_n b$', but that '$\exists_n S(a, b)$' does not.

As for the absolute quantifier 'exists_a', some lightweight realists may deny whether there is any such expression or any concept for it to express, though some may allow that there are such expressions while holding that they are hopelessly defective. Ontological anti-realists, as characterized here, hold that there are at least expressions that express the absolute quantifier, and that there may or may not be a concept of absolute quantification, but in any case the absolute quantifier is defective in such a way that there are no (or few) nontrivial truths that are expressed with it. Heavyweight realists, of course hold that there is an absolute quantifier and it is nondefective.

Remaining disagreements concern the relationship between these quantifiers and words such as 'exists'. Heavyweight realists and anti-realists[25] hold that in ontological contexts, 'exists' has the content of 'exists_a', while lightweight realists hold that it has the content of a lightweight quantifier such as 'exists_m' or 'exists_c'. All of these views can allow that in ordinary contexts, 'exists' has at least the correctness-conditions given by one or more lightweight quantifiers. Lightweight realists will also allow that it has these truth-conditions, as will many (but perhaps not all) ontological anti-realists, and some (but certainly not all) ontological realists. Those who allow this

Overall, there is certainly a substantive disagreement between heavyweight realists and the other two views, over whether there are nontrivial absolutely quantified truths. And there is certainly a semantic disagreement between lightweight realism and the other two views over the content of 'exists' in ontological contexts. There may also be a substantive disagreement between lightweight realists and ontological anti-realists about whether there is even a concept (or a possible expression) for absolute quantification, although this is less clear. In any case, it is once again clear that once we prescind from issues about language and concepts, lightweight realism and ontological anti-realism have more in common with each other than either has in common with heavyweight realism.

What about ontological relativism?

I have not yet discussed ontological relativism, on which ontological assertions have an assessment-relative truth-value. This view is closely connected to the ontological indeterminism I have outlined. Instead of holding that ontological existence assertions have an indeterminate truth-value, assessible by supervaluating across admissible frameworks, we could hold that they have an assessment-relative truth-value, assessible only by the standards of different frameworks. This

[25]There is also ontological contextualism, which holds that 'exists' has different contents in different *ontological* contexts, presumably expressing different lightweight quantifiers in each. This sort of ontological anti-realist, unlike mine, might deny that 'exists' ever expresses the absolute quantifier.

view would in some respects be parallel to the contextualism about ordinary existence assertions outlined earlier, except that in this case the framework serves as part of the context of assessment of an ontological existence assertion, not as part of the context of utterance of an ordinary existence assertion.

This package of views, combining contextualism about ordinary existence assertions with relativism about ontological existence assertions, has its attractions. The difference between the treatment of ordinary and ontological assertions is well-suited to reflect the different ways that disagreement functions in ordinary and ontological discourse, as outlined in sections 3 and 4. Ultimately, however, I think that relativism here is less well-motivated than in other domains. Ontological practice aims to prescind from the commitments of local frameworks and aims at objective truth, in a way that runs deeper than practice in domains where relativism has been attractive. I think that most participants in this practice would agree that *if* there is no framework-independent fact of the matter about ontology, then ontological assertions lack a determinate truth-value.

One can see relativism in other domains (aesthetics, predicates of personal taste, epistemic modality, perhaps morality) as an attempt to reconcile three features of discourse in these domains: (i) the sense that when two speakers utter S and $\neg S$, they are disagreeing, (ii) the lack of perspective-independent facts of the matter about who is correct in such cases, and (iii) the intuition that such utterances in these domains can be true or false, despite (i) and (ii). In the ontological domain, (i) and (ii) are present, but (iii) is less well-motivated. In particular, (ii) undercuts the intuition of truth or falsity more strongly here than elsewhere, in part because the presupposition of objectivity is much more central to ontological practice, and in part because ontological practice plays a much less central role in everyday discourse. So I think that ontological indeterminism is a better view. But ontological relativism is at least a close relative.

What about singular terms?

It has sometimes been objected to various sorts of ontological pluralism that such views cannot stop at giving multiple meanings to quantifiers. For example, Eklund (this volume) notes that the same sort of considerations apply to sentences containing singular terms. Say that 'Fred' is a term stipulated to refer to a cupcup on this table if it refers at all. Then insofar as 'Cupcups exist' is true on some readings but not on others, then 'Fred is on the table' will be true on some readings but not others. So the pluralist also needs to account for variability in the contents of singular terms.

In response: the current framework applies to singular terms as well as quantifiers. The frame-

work says that all expressions, not just quantified expressions, are associated with a function from furnished worlds to extensions. A singular term such as 'Fred' will have an extension in a furnished world that admits cupcups, and it will not have an extension in a furnished world that does not. So in contexts with a furnishing function that admits cupcups, 'Fred is on the table' may be true, and in contexts with a furnishing function that does not admit cupcups, 'Fred is on the table' will be false (or indeterminate, depending on the semantic treatment of empty terms).

Does this mean that even names such as 'Fred' are context-dependent? Yes and no. In different contexts, uses of the name will be will be associated with the same functions from furnished worlds to extensions, so these contents will not be context-dependent. But they will be associated with different functions from worlds to extensions, so these contents will be context-dependent. However, the source of the context-dependence will be the quite general variability of ontological frameworks, rather than anything specifically to do with names.[26]

In his paper in this volume, Eklund raises a related challenge for Carnapian varieties of ontological pluralism. Say that A is in a context with a restrictive framework that rejects cupcups, while B is in a context with a liberal framework that accepts cupcups. What should A say about the truth-value of an utterance S of 'Fred is on the table' by B? If A can truly say 'S is false', then we have some sort of relativism about utterance truth. And if A can truly say 'S is true', she must deny Tarski's thesis (here slightly modified from Eklund) that if an utterance of the form 'F(a)' is true, where 'a' is a singular term, there must be an entity that the utterance of 'a' refers to.

In response: In the current framework, this issue will arise only for assertions in ordinary contexts, not ontological contexts. In that case, I think A can and should truly say 'S is true', recognizing (as in section 3) that B has a different framework. As for Tarski's thesis, in the current account we have in effect modified it into the following: if an utterance of a sentence of the form 'F(a)' is true, then the utterance of 'a' must *quasi-refer*.[27] Here (speaking from within a liberal framework) we can say that an utterance of 'a' quasi-refers if it has an extension in the domain determined by the context and world of utterance. Of course if A is in a sufficiently impoverished context, she will not be able to use this characterization of quasi-reference, but that is just another

[26]Of course, if one holds that correctness is distinct from truth, then the truth-conditions of these sentences, and the corresponding application-conditions of names, need not be context-dependent. The context-dependence will only affect pragmatic correctness-conditions. Even if correctness is truth, on some proposals about the logical form of these sentences, the context-dependence may be traceable to variability in some other hidden element of a sentence (for example, a parameter for a framework), which combines with a context-independent name to yield context-dependence for the sentence as a whole.

illustration of the difficulty of doing semantics in a restricted framework.

Furthermore: if one wants to preserve Tarski's thesis, one must simply work within a maximally liberal framework, one that admits all the entities admitted by any framework. If one does, one can utter the original version of Tarski's thesis truly. In more restricted contexts, an utterance of the original version of Tarski's thesis will be false, just as in restricted contexts, an utterance of Euclid's theorem, 'There are an infinite number of primes', will be false. Here we find another example of the virtues of doing semantics in a liberal context.

Can we apply this framework to other ontological debates?

So far I have mainly applied this framework to debates over whether there are abstract objects, macroscopic objects, and mereological sums. The framework applies fairly straightforwardly to any ontological debate over whether enities of a certain sort exist. What about debates that do not fit quite so straightforwardly? For example, the debate over coincidence: are the statue and the lump of clay one thing or two? The debate over persistence: are objects three-dimensional objects that persist by enduring, or four-dimensional objects that persist by perduring? The debate over presentism: do the past and the future exist, or just the present?

I think the framework applies straightforwardly to the debate over coincidence. There are furnishing functions that never supply coincident objects. For example, on one furnishing function, everything that exists is a sum of simple parts, and has precisely those simple parts whenever it exists. That furnishing function delivers a furnished world without coincident objects. It is plausible that in that furnished world, our terms 'that lump' and 'that statue' pick out the same sum of particles: a sum that can persist even once it is no longer shaped as a statue, but not after some of the particles are destroyed. Other furnishing functions supply many coincident objects. A maximally liberal furnishing function might supply objects objects for any set of persistence conditions (both temporal and modal), delivering a furnished worlds with many coincident objects. It is plausible that in that furnished world, our terms 'that lump' and 'that statue' pick out two different but coincident objects with different persistence conditions: one that can survive the loss of particles but not the loss of a statue shape, while another for which it is the other way around.

[27] In response to a move that has something in common with this, Eklund suggests that the argument could work using a modified thesis holding that if 'a' is an ontologically committing term, 'F(a)' is true iff 'a' refers, where this sentence is more or less true by definition of "committing". But then A will deny that B's term 'a' is committing. On the current picture, singular terms used in ontological contexts may be committing in this sense, but singular terms used in ordinary contexts are not.

I am inclined to think that both furnishing functions are admissible, so that there is no fact of the matter as to which ontological view is correct.

Something similar may apply to the debate over perdurance versus endurance, depending on exactly how that debate is formulated. One might in principle apply the framework to the debate over presentism, with furnishing functions that admit past and future objects, or just present objects. However, in this case it is far from obvious that both furnishing functions will be admissible, because of constraints imposed by the fundamental level. For example, it seems plausible that fundamental truths will either make reference to past and future microphysical states of the world, or they will not. Either way, the debate over presentism may arise already at that level. This coheres with the general picture suggested earlier, on which the most substantive metaphysical debates are those that arise on the fundamental level.

12 Conclusion

It will be obvious that my picture is closely parallel to Carnap's. Ordinary and ontological existence assertions correspond to internal and external questions. Ordinary existence assertion, like internal questions, have truth-conditions that are determined by an associated framework, and have truth-values that are determined trivially by the framework and the underlying state of the world. Ontological existence assertions, like external questions, purport to be independent of a framework, but lack determinate truth-conditional content, and typically lack determinate truth-values. Ontological discourse is thereby often defective, but ordinary discourse about existence is not defective at all. Of course there are differences between my approach and Carnap's. But I hope that the present treatment suggests that a Carnapian approach to ontology is coherent and viable.

Bibliography

Bennett, K. (this volume). Composition, coincidence, and metaontology. In Chalmers, Manley, and Wasserman, forthcoming.

Carnap, R. 1950. Empiricism, semantics, and ontology. *Revue Internationale de Philosophie* 4: 20-40. Reprinted as an appendix to *Meaning and Necessity: A Study in Semantics and Modal Logic*. University of Chicago Press.

Chalmers, D.J. 2004. Epistemic two-dimensional semantics. *Philosophical Studies*.

Chalmers, D.J. forthcoming. Terminological disputes. See http://consc.net/papers/terminology.ppt.

Chalmers, D.J., Manley, D. & Wasserman, R. forthcoming. *Metametaphysics: New Essays on the Foundations of Ontology*. Oxford University Press.

Dorr, C. 2005. What we disagree about when we disagree about ontology. In (Mark Kalderon, ed.) *Fictionalist Approaches to Metaphysics*. Oxford University Press.

Egan, A., Hawthorne, J. & Weatherson, B. 2005. Epistemic modals in context. In (G. Preyer and G. Peter, eds.) *Contextualism in Philosophy*. Oxford University Press.

Eklund, M. (unpublished) Neo-Fregean ontology. http://www.people.cornell.edu/pages/me72/pnfo.pdf

Eklund, M. (this volume). Carnap and logical pluralism. In Chalmers, Manley, and Wasserman forthcoming.

Field, H. 1980. *Science Without Numbers*. Princeton University Press.

Field, H. 1993. The conceptual contingency of mathematical objects. Mind 102: 285-299.

Hale, B. & Wright, C. 1992. Nominalism and the contingency of abstract objects. *Journal of Philosophy* 89:111-135.

Hale, B. & Wright, C. 2001. *The Reason's Proper Study*. Oxford University Press.

Hirsch, E. 1993. *Dividing Reality*.

Hirsch, E. 2002. Quantifier variance and realism. *Philosophical Issues* 12: 51-73.

Hirsch, E. 2005. Physical-object ontology, verbal disputes, and common sense. *Philosophy and Phenomenological Research* 70:67-97.

Hofweber, T. 2005. A puzzle about ontology. *Nous* 39:256-283.

Horgan, T. 2001. Contextual semantics and metaphysical realism: Truth as indirect correspondence. In (M. Lynch, ed.), *The Nature of Truth*. MIT Press.

Horgan, T. & Potrc, M. 2006. Abundant truth in an austere world. In (P. Greenough and M. Lynch, eds) *Truth and Realism*. Oxford University Press.

Korman, D. forthcoming. Unrestricted composition and the argument from vagueness.

Koslicki, K. 2004. The crooked path from vagueness to four-dimensionalism. *Philosophical Studies* 114:107-34.

Macfarlane, J. 2005. Making sense of relative truth. *Proceedings of the Aristotelian Society* 105: 321-39.

Price, H. (this volume). The ghost who walks: Did metaphysics survive the Quine-Carnap debate? In Chalmers, Manley, and Wasserman forthcoming.

Putnam, H. 1987. *The Many Faces of Realism*. Open Court.

Putnam, H. 2004. *Ethics without Ontology*. Harvard University Press.

Quine, W.V. 1948. On what there is. Review of Metaphysics 2:21-38.

Quine, W.V. 1951. On Carnap's views on ontology, Philosophical Studies 2:65-72.
Rayo, A. forthcoming. Ontological commitment.
Rayo, A. & Uzquiano, G. (eds) forthcoming. *Absolute Generality*. Oxford University Press.
Schaffer, J. 2003. Is there a fundamental level? *Nous* 37:498-517.
Sidelle, A. 2002. Is there a true metaphysics of material objects? Philosophical Issues 12:118-145.
Sider, T. 2001. *Four-Dimensionalism*. Oxford University Press.
Sider, T. (this volume). Ontological Realism. In Chalmers, Manley, and Wasserman, forthcoming.
Stanley, J. 2005. *Knowledge and Practical Interests*. Oxford University Press.
Thomasson, A. (this volume). Answerable and unanswerable questions. In Chalmers, Manley, and Wasserman forthcoming.
van Inwagen, P. 1990. *Material Beings*. Ithaca: Cornell University Press.
van Inwagen, P. 1998. Meta-ontology, Erkenntnis 48: 233-250.
van Inwagen, P. 2002. The number of things. Philosophical Issues 12:176-196.
van Inwagen, P. (this volume). In Chalmers, Manley, and Wasserman, forthcoming.
Yablo, S. 2000. A paradox of existence. In (A. Everett and T. Hofweber, eds.) *Empty Names, Fiction and the Puzzles of Non-Existence*. University of Chicago Press.