# Meta-Ontological Minimalism

Øystein Linnebo
University of Bristol

Draft of March 13, 2008; slightly edited June 26, 2008

## 1 Introduction

Kant famously argued that all existence claims are synthetic.[1] An existence claim can never be established by conceptual analysis alone but will always require some appeal to intuition or perception, thus making the claim synthetic. This view is boldly rejected in Frege's *Foundations of Arithmetic* ((Frege, 1953)), where Frege defends an account of arithmetic which combines both platonism and logicism. Frege's platonism consists in taking arithmetic to be about real and independently existing objects. And his logicism consists in taking the truths of pure arithmetic to rest on just logic and definitions and thus be analytic. Most philosophers now probably agree with Kant in this debate: the existence claims of Frege's platonism cannot be established on the basis of logic and conceptual analysis alone. This is why George Boolos, only slightly tongue-in-cheek, can offer a one-line refutation of Fregean logicism: "Arithmetic implies that there are two distinct numbers" ((Boolos, 1997), p. 302), whereas logic and conceptual analysis—Boolos takes us all to know—cannot underwrite any existence claims (other than perhaps of one object, so as to streamline logical theory).

However, the disagreement between Kant and Frege is alive and well in a somewhat different form. Forget the problematic analytic-synthetic distinction. Can there be objects which are "thin" in the sense that very little is required for their existence? A classic example of an appeal to thin objects is the view in the philosophy of mathematics that the mere consistency or coherence of a mathematical theory suffices for the existence of the objects that the theory purports to describe. This view has been held by many leading mathematicians

[1] See (Kant, 1997), esp. B622-3.

and continues to exert a strong influence on contemporary philosophers of mathematics.[2] A more recent example is the neo-Fregean view that the equinumerosity of two concepts is conceptually sufficient for the existence of the number that specifies the cardinality of both concepts. For instance, the fact that the knives on the table can be one-to-one correlated with the forks on the table is said to be conceptually sufficient for the existence of a number that specifies the cardinality both of the knives and of the forks.[3]

In these examples certain objects are said to be thin in an *absolute* sense. But objects can also be regarded as thin in a *relative* sense. Say that an object $x$ is thin relative to some other objects if, given the existence of these other objects, very little is required for the existence of $x$. Someone attracted to the view that pure sets are thin in the absolute sense is then likely also to be attracted to the view that an impure set is thin relative to any urelemente (or non-sets) that enter into the set. On this view the existence of a set of all the books in my office would require very little over and above the existence of the books. Another example is the view that a mereological sum is thin relative to its parts. On this view the existence of a mereological sum of all my books would require very little over and above the existence of these books.[4]

If the defenders of thin objects (in either the absolute or relative sense) are regarded as heirs to the Fregean view that there are analytic existence claims, then there are also lots of heirs to the contrasting Kantian view. For instance, Hartry Field has attacked views according to which mathematical objects are thin, sometimes mentioning the Kantian origin of his criticism.[5] And various metaphysicians have attacked the view that mereological sums are thin relative to their parts.[6] Like the original Kantian view on analytic existence claims, the contemporary view that there can be no thin objects strikes many people as highly plausible. Appeals to thin objects often come across as attempts to pull rabbits out of a hat.

The goal of this article is to explain, defend, and demystify the idea of thin objects. I will refer to the view that there are thin objects as *meta-ontological minimalism* or sometimes just *minimalism* for short. Let me explain the label. Ontology is of course the study of what there is. Meta-ontology, on the other hand, is the study of these key concepts of ontology,

[2]See for instance (Parsons, 1990), (Resnik, 1997), and (Shapiro, 1997).
[3]See for instance (Wright, 1983) and the essays collected in (Hale and Wright, 2001).
[4]For examples of philosophers attracted to this view see (Lewis, 1991), esp. Section 3.6 and (Sider, 2007).
[5]See (Field, 1989), pp. 5 and 79-80.
[6]See for instance (Rosen and Dorr, 2002) [better examples?].

such as the concepts of existence and objecthood. Meta-ontological minimalism is accordingly the view that the key concepts of ontology have a minimal character. Not surprisingly, this view tends to result in very generous ontologies. For the less that is required for existence, the more objects there will be.[7]

However, it is important to note that meta-ontological minimalists do not claim that *all* objects are thin and that their existence thus only makes some minimal demand on reality. The minimalist claim is that the notion of an object itself is thin and thus *allows for* thin objects. But they happily admit that additional thickness may come from the kind of object in question. For instance, most minimalists think that elementary particles are thick and that their existence thus makes some substantive demand on reality. But the minimalists insist that the thickness of elementary particles derives from what it is to be an elementary particle, not from what it is to be an object.

This paper is structured as follows. First I review some arguments in favor of thin objects and describe three examples of views that posit such objects. Next I outline some problems that these views face. I then turn to my main goal, which is to develop a version of meta-ontological minimalism and a conception of thin objects. This view is based on a Fregean notion of abstraction. But my view differs in important ways from the neo-Fregean view of Bob Hale and Crispin Wright; in particular, I tie the notion of an object to that of a semantic value and make crucial use of a principle of compositionality for semantic values. I end by arguing that my version of minimalism is well equipped to solve the problems that minimalist views face.

## 2 The appeal of thin objects

Meta-ontological minimalism appears to enjoy a number of appealing features. Perhaps most importantly, it promises a way to accept face value readings of discourses whose ontologies would otherwise be philosophically problematic.

Arithmetic provides a good example. The language of arithmetic contains a variety of proper names which (it seems) are supposed to refer to certain abstract objects, namely the natural numbers. The language also contains quantifier phrases which (it seems) are supposed

[7]Indeed, Matti Eklund labels as 'maximalism' the kind of ontological view encouraged by one version of what I call meta-ontological minimalism. See (Eklund, 2006) and (Eklund, ).

to range over the natural numbers. Moreoever, a great variety of theorems expressed in this language appear to be true. For a lot of such theorems are asserted in full earnest by educated lay people as well as professional mathematicians. And since the arithmetical competence of these people is beyond question, there is reason to believe that most of their arithmetical assertions are true. But if these theorems are true, then their various subexpressions must succeed in doing what they are supposed to do. In particular, their singular terms and quantifiers must succeed in referring to and ranging over natural numbers. And for this kind of success to be possible, there must exist abstract mathematical objects.

This is a powerful argument. But is it sound? Since all I did was to observe that the premises *appear* to be true, they can of course be challenged. However, it is *prima facie* attractive to take these appearances at face value. In particular, this will save us the difficult task of having to show how both lay people and experts can be deceived about something they take to be obviously true. And when the premises are taken at face value, the argument shows that there must exist abstract mathematical objects.

However, this ontology of abstract objects is often found to be philosophically problematic. One well-known worry concerns epistemic "access" to such objects. Since perception and all forms of instrumental detection are based on causal processes, these methods cannot give us access to abstract objects such as the natural numbers. How then can we acquire knowledge of them?[8] Another worry is the sheer extravagance of postulating such huge ontologies. How can we postulate an infinity of new objects with such a light heart? No physicist would so unscrupulously postulate an infinity of new physical objects. Why then should mathematicians get away with it? Philosophers are of course divided about how serious these worries are. But any successful account of mathematical objects needs to have some response to the worries, even if only to explain why they are misguided.

Thin objects offer an extremely promising strategy for responding to the worries. The vast ontology of mathematics may well be problematic when it is understood in a thick sense. If mathematical objects were pretty much like elementary particles except for being abstract, then there would indeed be good reason to worry about epistemic access and ontological extravagance. But perhaps mathematical objects need not be understood in this way. If there are such things as thin objects, then the existence of mathematical objects need not make

[8] This worry was made famous by (Benacerraf, 1973). For discussion and improvements, see (Field, 1989) and (Linnebo, 2006).

much of a demand on the world. It may for instance suffice that the theory purporting to describe the relevant mathematical objects is coherent. And although facts about the coherence of mathematical theories are still inadequately understood, they are less problematic than thick mathematical objects would be. It is at least not a complete mystery how we can have epistemic access to facts about the coherence of mathematical theories. And since this account of mathematical objects sets the bar to existence extremely low, it is not at all surprising that an extravagant ontology should result.

Arithmetic is only one example of how thin objects can be philosophically useful. Other examples abound. Consider for instance the philosophical debate about the existence of mereological sums. Very often we speak as if there are various kinds of mereological sums; for instance, we speak as if there are decks of cards, bunches of grapes, crowds of people. And many of these claims appear to be true. But some philosophers find mereological sums to be problematic, often because of worries similar to the ones discussed above. If mereological sums are thin relative to their parts—that is, if little or nothing is required for their existence other than the existence of their parts—then we would be in a good position to address these worries.

These explanations will obviously have to be spelled out further to be fully convincing. But this explanatory strategy is promising enough to make the idea of thin objects extremely attractive. If some form of the idea can be articulated and defended, its explanatory potential may be enormous.

## 3 Three approaches to thin objects

I now outline three approaches to thin objects.

### 3.1 A coherentist approach

I begin with the view that the coherence of a mathematical theory suffices for the existence of the objects that the theory purports to describe. For instance, since it is coherent to supplement the ordinary real number line $\mathbb{R}$ with two infinite numbers $-\infty$ and $+\infty$, the so-called *extended real number line* $\bar{\mathbb{R}} = \mathbb{R} \cup \{-\infty, +\infty\}$ exists. And since it is coherent to supplement $\mathbb{R}$ with the imaginary unit $i = \sqrt{-1}$ and other (non-real) complex numbers, the complex number field $\mathbb{C}$ exists. All that the existence of these new mathematical objects amounts to is, according to this view, the coherence of the theories that describe the relevant structures. I will therefore refer to this as a *coherentist* approach to thin objects.

This approach enjoys wide-spread support from within mathematics itself, where it has been defended by a number of very prominent mathematicians. For instance, in his famous correspondence with Frege, David Hilbert wrote:

> As long as I have been thinking, writing and lecturing on these things, I have been saying the exact reverse: if the arbitrarily given axioms do no contradict each other with all their consequences, then they are true and the things defined by them exist. This is for me the criterion of truth and existence.[9]

A similar view is voiced by Georg Cantor, whose general philosophical outlook is very different from Hilbert's:

> Mathematics is in its development entirely free and only bound in the self-evident respect that its concepts must both be consistent with each other and also stand in exact relationships, ordered by definitions, to those concepts which have previously been introduced and are already at hand and established.[10]

In philosophy too the coherentist approach to thin objects has enjoyed wide-spread support, especially in the past few decades. Its defenders include central contemporary philosophers of mathematics such as Charles Parsons, Michael Resnik, and Stewart Shapiro.[11]

### 3.2 An abstractionist approach

A second approach to thin objects is found in the neo-Fregean philosophy of mathematics developed by Hale and Wright. The neo-Fregeans seek to provide a logical and philosophical foundation for classical mathematics on the basis of so-called *abstraction principles.* These are principles of the form

$$(*) \qquad \Sigma(\alpha) = \Sigma(\beta) \leftrightarrow \alpha \sim \beta$$

[9]Letter to Frege of 29 December 1899, in (Frege, 1980).
[10]See Cantor 1883 ((Ewald, 1996), p. 896). For two examples from Henri Poincaré and another from Hilbert, see (Ewald, 1996), pp. 1026, 1055, and 1105.
[11]See for instance (Parsons, 1990), (Resnik, 1997), and (Shapiro, 1997).

where $\alpha$ and $\beta$ range over items of some sort, where $\sim$ is an equivalence relation on such items, and where $\Sigma$ an operator that maps such items to objects. The neo-Fregeans are particularly fond of *Hume's Principle*, which says that the number of $F$s (symbolized as $\#F$) is identical to the number of $G$ just in case the $F$ and the $G$s can be one-to-one correlated (symbolized as $F \approx G$):

(HP) $$\#F = \#G \leftrightarrow F \approx G$$

This principle has the amazing mathematical property that, when added to second-order logic along with some definitions, we are able to derive all of ordinary (second-order Peano-Dedekind) arithmetic.[12] Abstraction principles are available for many other kinds of abstract object as well, for instance directions, geometrical shapes, and linguistic types. Some further examples will be developed below.

If true, an abstraction principle will provide unproblematic epistemic and semantic access to the objects invoked on the left-hand side of an abstraction principle: we simply proceed via its unproblematic right-hand side. However, what reason do nominalists (and other non-believers) have to accept Hume's Principle and other abstraction principles as true? The neo-Fregean response to this challenge turns on regarding the objects invoked on the left-hand side as thin.

This response can be developed in the form of a view of what is required for a singular term $t$ to refer.[13] At the very least the term must have sense. But more is presumably required. What is this further requirement? The question can be put in terms of the following equation:

(E) $$t \text{ has sense} + X \;\Leftrightarrow\; t \text{ has reference}$$

where '$\Leftrightarrow$' means something like mutual conceptual entailment. That is, what requirement $X$ do we have to add to the claim that $t$ has sense to get something that is conceptually equivalent to the claim that $t$ has reference?

When $t$ is an *abstraction term*—that is, a term of the form '$\Sigma(\alpha)$'—then Hale and Wright claim that the further requirement $X$ that is needed to advance from sense to reference is

[12] This result, which is known as *Frege's Theorem*, is hinted at in (Parsons, 1965) and explicitly stated and discussed in (Wright, 1983). For a nice proof, see (Boolos, 1990).

[13] See (Hale and Wright, ).

just that the item $\alpha$ associated with $t$ be equivalent to itself: $\alpha \sim \alpha$. That is, they propose the following abstractionist solution to the equation (E):[14]

$$\text{(A)} \qquad \text{`}\Sigma(\alpha)\text{' has sense} + (\alpha \sim \alpha) \;\Leftrightarrow\; \text{`}\Sigma(\alpha)\text{' has reference}$$

According to this view, the left-hand side of (A) conceptually entails the right-hand side and its claim that the abstraction term '$\Sigma(\alpha)$' refers. The notions of reference and objecthood have been "scaled" so as to ensure that (A) comes out right.[15] This abstractionist approach to thin objects will serve as the starting point for the approach to be developed in Section 7.

### 3.3 A priority-based approach

A third approach to thin objects has been articulated in recent work by Matti Eklund. This approach is based on the idea that truth is somehow "prior to" reference. Consider an atomic sentence $S$ in which an expression $t$ functions as a singular term. Then the truth of $S$ is "constitutively prior" to the reference of $t$; in particular, the question of $S$'s truth can be settled prior to and independently of the question of $t$'s reference.[16] For instance, let $S$ be an atomic sentence in the language of pure arithmetic. Then the truth of $S$ should be assessed by ordinary arithmetical criteria. If these criteria reveal $S$ to be true, we may conclude that the singular term $t$ refers to a natural number. The correct approach is thus to begin with the question of truth and on this basis establish a fact about reference. The nominalist's mistake is according to this view to adopt the reverse approach. The nominalist insists on a prior and independent assurance that $t$ refers. Only on this basis is the nominalist willing to pronounce on the truth or falsity of $S$. But this reverses the natural priority of truth to reference.

Eklund claims to find this view—which I will call *priority-based minimalism*—in the work of Dummett, Hale, and Wright.[17] He also argues that the view is of considerable independent interest but stops short of actually endorsing it.

It is not hard to see that priority-based minimalism leads to a very generous ontology. To show that there are objects of some possible kind $K$, it suffices to find a sentence that

[14] Hale and Wright do not say whether this solution is unique. So although they defend an abstractionist form of minimalism, nothing they say rules out there being other forms as well.

[15] See (Hale and Wright, 2008), p. 26 for this metaphor of "scaling."

[16] The characterization "constitutively prior" is used at (Eklund, 2006), p. 100.

[17] His strongest case for this exegetical claim seems to me to be the response to nominalism that is developed in (Wright, 1992). Another strong case is (Dummett, 1956).

purports to refer to $K$s and which is such that the only obstacle to its truth would be the failure of some of its singular terms to refer. If the question of truth can be settled prior to and independently of the question of reference, then this sentence will be guaranteed to be true and its singular terms will accordingly refer to $K$s.[18]

## 4 Problems with thin objects

Although the idea of thin objects holds great promise, it also faces a number of problems and challenges. I now outline some of the more important ones, which any successful version of meta-ontological minimalism must be capable of addressing.

### 4.1 The problem of existence

Why should we believe that there are such things as thin objects in the first place? Consider the issue from the point of view of a nominalist. The minimalist claims that there are certain objects whose existence amounts to nothing more than the obtaining of certain apparently innocent facts. These apparently innocent facts can for instance be the coherence of a mathematical theory or the equinumerosity of a concept with itself. But why should the nominalist accept that the obtaining of such facts suffices for the existence of certain controversial objects that play no apparent role in these facts? What prevents her from accepting these facts while still denying that this suffices for the existence of any of the controversial objects?

This response from the nominalist will no doubt strike a lot of philosophers as perfectly reasonable. And as mentioned above, the nominalist can here also appeal to the familiar Kantian doctrine that conceptual analysis alone can never underwrite the truth of any existence claim.

### 4.2 The problem of overgeneration

Assume that the problem of existence can be solved. Then the question arises how thin various kinds of object actually are. Few minimalists would want to claim that elementary particles are thin. For the existence of such objects seems to make a substantive demand on the world. But once minimalists have allowed some kinds of object to be thin, what

[18]There are many other approaches to thin objects as well, which cannot be discussed here. [See for instance (Balaguer, 1998) (although this may be a clearer case of ontological maximalism than of meta-ontological minimalism), Tait? Linsky & Zalta?].

right do they have to deny that elementary particles are thin as well?[19] Minimalists owe us an explanation of what it is that makes abstract objects thin and elementary particles comparatively thick. Consider for instance the view that the coherence of a theory of pure mathematics suffices for the existence of the mathematical objects described by the theory. Advocates of this view owe us an explanation of why the same is not true of particle physics. Without some such explanation the various minimalist views will overgenerate and result in vast numbers of undesired objects.

Worries about overgeneration arise in the abstract domain as well. Consider a theory that is consistent but utterly arbitrary and unintuitive. Do there really exist mathematical objects of the sort described by the theory? Or would such objects exist provided the theory met some further requirement, such as categoricity?[20]

### 4.3 The problem of lack of uniformity

Assume that the problem of overgeneration can be solved. Then minimalists still face the question how much thick and think objects have in common. Given the dramatic differences that exist between thick and thin objects, one wonders whether the two kinds of item belong under the same rubric at all. Perhaps the words 'existence' and 'object' are being used ambiguously. If so, it might be better not to apply the label 'object' indiscriminately to both kinds of item but rather to use different labels altogether.

Analogous questions arise concerning the notion of reference. Are cases of reference to thick and thin objects really instances on one and the same common phenomenon? Dummett argues that they are not and that we instead are dealing with two distinct phenomena. He claims that reference to mathematical objects is "semantically idle" in that the "notion of the reference of the term, as determined by its sense, plays no role in our conception of what determines the thought as true or false" ((Dummett, 1991), p. 193).[21]

### 4.4 The problem of consistency

The last and potentially most fatal problem is the threat of inconsistency. The threat is most dramatic in the case of the priority-based approach. Eklund argues that, at least in mathe-

[19] Or, taking a cue from (Field, 1984), what is prevent a theist from invoking the idea of thin objects to defend some easy route to the existence of God?

[20] A theory is said to be *categorical* if it has a unique model up to isomorphism.

[21] See for instance (Dummett, 1981), pp. 498-511 and (Dummett, 1991), chapters 15-16.

matical cases, the possible existence of objects of some kind $K$ suffices for the actual existence of $K$s. But as he observes, inconsistency then looms: for there are kinds of mathematical objects which individually appear to be possible but which provably cannot co-exist. George Boolos offers a nice example. The numbers described by Hume's Principle can only exist in an infinite universe. But there is another abstraction principle, known as *the nuisance principle*, which is satisfiable only in finite universes.[22] So although the two principles are individually consistent and the objects that each describes thus appear to be possible, the two principles cannot both be true, as there is no universe in which both are satisfiable.

A related worry arises for the abstractionist approach. For the acceptable abstraction principles turn out to be surrounded by unacceptable ones, which are incoherent or downright inconsistent, or which conflict with one another.[23]

The coherence-based approach appears to be better equipped to deal with the threat of inconsistency. For what this approach claims isn't that coherence guarantees *truth* but that it guarantees *satisfiability*. That is, the claim is that when a theory $T$ is coherent, there is a mathematical structure that satisfies $T$. This enables a nice response to the problem of individually coherent yet jointly incompatible theories: This approach is committed to each theory's being satisfied in some structure but not to both theories' being true in the actual universe.[24]

## 5 Objects as referents of singular terms

I will now develop a version of meta-ontological minimalism which I believe provides answers to all of the above problems. This version is inspired by the abstractionist approach described above. My argument proceeds in three steps. The first step glosses the notion of object as a possible referent of a singular term. The second step glosses the notion of a referent of a singular term in terms of the notion of a semantic value. The third step offers a minimalist account of what it takes for a singular term to possess a semantic value. Taken together, the three steps provide a minimalist account of what is required for the existence of an object.

[22] Consider the equivalence relation that holds between two concepts $F$ and $G$ just in case they are coextensive apart from at most finitely many exceptions. The associated abstraction principle can be shown to be satisfiable in all and only finite domains. See (Boolos, 1990).

[23] This is known as *the bad company problem*. See (Linnebo, b) for an introduction and further references.

[24] But see (Hellman, 2001), pp. 191-2 for worries about the consistency of one of the more developed ways of spelling out this approach, namely Stewart Shapiro's.

Let's begin with the first step. One may wonder how an inquiry into the concept of an object can even be possible. For as Frege observes, this concept is "too simple to admit of logical analysis," and "what is logically simple cannot have a proper definition."[25] But although Frege is no doubt right that a "proper definition" of the concept of an object is out of the question, I believe it is both possible and reasonable to ask for some further explication of the concept. Even if the concept cannot be *defined* in more basic terms, it can still be glossed or characterized, for instance by relating it to other concepts and by explaining the role that it plays in our thought and reasoning. It is instructive to compare with the notion of conjunction. Although this notion too is a primitive which cannot be defined in more basic terms, a lot can be said to gloss or characterize it. We can for instance describe its inferential properties and its possession conditions.

So let's examine the role that objects play in our semantic theories.[26] There appear to be two different but related roles: objects serve as referents of singular terms and as values of bound first-order variables. Frege's explication of the notion of object focusses on the former role of objects as the referents of singular terms. He takes objects to be the kinds of item that singular terms refer to. Quine, on the other hand, focuses on the latter role of objects as values of bound variables, as encapsulated in his famous slogan that "to be is to be the value of a bound variable."[27] Since the referent of any singular term can also serve as the value of a bound variable, it follows that everything that is an object in Frege's sense is also an object in Quine's sense. I will return to the question of whether the converse holds.

Which explication is better? My own preference is for Frege's explication over Quine's. For I take singular terms and their reference to be more fundamental than quantifiers and their ranges.[28] Quantification is explained in terms of its relation to its instances, which involve singular terms. I therefore propose the following Fregean alternative to Quine's slogan: To be is to be a possible referent of a singular term.

It may be objected that this Fregean explication makes the notion of an object excessively semantic. As an initial defense, I observe that the proposed Fregean explication is no more semantic than the more familiar Quinean explication. Where the former talks of the reference

[25]See respectively (Frege, 1891), p. 140 and (Frege, 1892), p. 182.

[26]A similar examination is found in (Dummett, 1981), ch. 14.

[27]By "bound variable" Quine here means bound first-order variable, as these are the only variables he thinks can legitimately be bound. See (Quine, 1986). But the Quinean gloss on the notion of an object is independent of this: we just need to add that the slogan is concerned with bound *first-order* variables.

[28]See also [Dummett and Wright. Also Hale?].

of singular terms, the latter talks of the ranges of quantifiers. The only significant difference is the use of modal vocabulary in the Fregean explication: an object is a *possible* referent of a singular term. This is required because not every object is the referent of some actual singular term, whereas every object is in the range of some actual quantifier.[29] But it is hard to see why this difference should make the Fregean explication more vulnerable to the present objection than the Quinean explication.

A better strategy for the objector is to argue that *both* explications are unacceptably semantic. Then a more thoroughgoing defense will be needed. In particular, it will then need to be shown that the use of semantic notions to gloss the notion of an object does not make objects or their existence depend on there being a language. I believe this can be shown quite convincingly. Both Frege and Quine exploit the contingent fact that we have a language to gloss one of our key ontological notions. But nothing prevents the notions thus glossed from being applied to other possible worlds, including ones at which there is no language. For when we entertain the counterfactual possibility that there is no language, we are still allowed to rely on the contingent fact that *we*—the people who evaluate the counterfactual—have a language.

Another objection to the Fregean explication of the notion of an object is that it is insufficiently general. We saw above that any object in the Fregean sense is also an object in the Quinean sense. But arguably the converse does not hold. The apparent failure of the converse is nicely brought out by a problem that has recently received a lot of attention.[30] The problem concerns mathematical structures with non-trivial automorphisms, that is, where objects can be permuted in a way that leaves the structure intact. A nice example are the imaginary units $i$ and $-i$ in complex analysis. These units are structurally indiscernible in the sense that every structural property that holds of one also holds of the other. This makes it hard to see how a singular term could determinately refer to one of these units rather than the other. Since the units are structurally indiscernible, nothing internal to the structure could ensure that the term refers to one rather than the other. And it is hard to see how facts external to the structure—whether mathematical or physical—could be responsible for a determinate reference relation of the sort in question. The upshot is that the imaginary

[29] At least this will be so if we can quantify over absolutely everything. But this assumption has been challenged. See (Rayo and Uzquiano, 2006), especially the introduction, for an overview of the ensuing discussion.

[30] See (Keränen, 2001), where the problem is used as an objection to non-eliminative forms of mathematical structuralism.

units appear not to qualify as objects in the Fregean sense. But it still seems that we can quantify over all the complex numbers, including the imaginary units. It thus seems that the complex numbers qualify as objects in the Quinean sense but not in the Fregean sense. And this may be a reason to prefer the Quinean explication over the Fregean one.

I see two possible responses to this objection. An *uncompromising response* would be to deny that we can make sense of quantification over objects to which we cannot determinately refer. For the standard truth-conditions for quantified formulas is based on the successive assignment of each object in the domain to the variable bound by the quantifier. But it is doubtful that this makes sense if the domain contains objects to which determinate singular reference is impossible.

A *compromising response* would be to point out that my primary concern in this section and the next is to argue that the notion of an object should be explicated in terms of the notion of a semantic value. It is not essential to this argument that the semantic values in question can be ascribed to individual singular terms and not only serve as values of bound variables. It is thus not essential to my project that we adopt the Fregean rather than the Quinean explication on the notion of an object. Defenders of the Quinean explication could probably develop and defend analogues of the next two steps of my argument. This compromising response has the advantage of allowing me to focus on my main opponent, namely people who object to my minimalism in both its Fregean and its Quinean form.

## 6 Reference as possession of semantic value

The second step of my argument relates the notion of a referent of a singular term to that of a semantic value. In order to explain this I first need to review some ideas from semantics and the philosophy of language. In these fields it is widely assumed that each component of a complex expression makes some definite contribution to the meaning of the complex expression. This contribution is known as its *semantic value*. I will write $[\![\mathbf{E}]\!]$ for the semantic value of an expression **E**. For instance, Frege held that the semantic value of a sentence is its truth-value and that the semantic values of other expressions are their contributions to the truth-values of sentences in which they occur. In particular, the semantic values of singular terms are just their referents.

It is also widely assumed that the meaning of a complex expression is functionally determined by the semantic values of its components and their syntactic mode of combination. This assumption is known as *compositionality.* For instance, according to Frege the semantic value of a simple sentence such as 'John runs' is determined by the equation:

$$\llbracket \text{John runs} \rrbracket = \llbracket \text{runs} \rrbracket(\llbracket \text{John} \rrbracket) \tag{1}$$

That is, the semantic value of the sentence 'John runs' is the result of applying the function which is the semantic value of the predicate 'runs' to the argument which is the semantic value of the subject 'John'. More generally, let $C$ be some syntactic operation applicable to syntactic expressions $E_1, \ldots, E_n$. Then there is some semantic operation $C^*$ corresponding to $C$ such that the semantic value of the result of applying the syntactic operation $C$ to the expressions $E_1, \ldots, E_n$ is identical to the result of applying the semantic operation $C^*$ to the expressions' semantic values:

$$\llbracket C(E_1, \ldots, E_n) \rrbracket = C^*(\llbracket E_1 \rrbracket, \ldots, \llbracket E_n \rrbracket). \tag{2}$$

Why should the ordinary notion of reference be explicated in terms of the technical notion of semantic value? One reason is that the notion of a semantic value generalizes the ordinary notion of reference. Where the notion of reference is intended primarily for singular terms, the notion of semantic value is meant to be applicable to expressions of all grammatical categories. Another reason is that the notion of semantic value carries with it much less intuitive baggage than that of a referent. The ordinary notion of a referent is naturally understood in a "thick" way. For instance, a referent is naturally taken to be something that one can somehow encounter, that plays an ineliminable role in the truth of predications, and that is completely independent of us and our representational devices. The technical notion of a semantic value carries no such baggage. Some of the respects in which the technical notion of a semantic value is "thinner" than the intuitive and pre-theoretic notion of a referent will be discussed in Section 9, where I will also argue that it is advantageous to set aside any "thick" connotations of the pre-theoretic notion of a referent. However, provided that these connotations are set aside, I have no objection to continued talk about the semantic values of singular terms as their referents.

# 7 When does a singular term have a semantic value?

The third step of my argument consists of a sufficient condition for a singular term to have a semantic value. Since this is the most important and distinctive step of the argument, it needs to be spelled out with some care.

Let's begin by getting clearer about the nature of this third step. The question is under what conditions a singular term has a semantic value. Let $\mathrm{SV}(t, a)$ be the relation that holds between a singular term $t$ and its semantic value $a$. For instance, $t$ can be an inscription of 'Tony Blair', and $a$, the recent British prime minister. What makes it the case that the former has the latter as its semantic value? Since the relation can hardly be a primitive one, there must be something that is responsible for its obtaining. Compare the relation of ownership, which also isn't a primitive one. So when I bear the ownership relation to my bicycle, there must be something responsible for the obtaining of this relation. The study of what it is in virtue of which expressions have semantic values is sometimes called *meta-semantics*.[31] One of the principal tasks of meta-semantics is thus to provide an account of what relations of the form $\mathrm{SV}(t, a)$ consist in.

## 7.1 The abstractionist answer restated

The sufficient condition for possession of semantic value that I wish to defend is inspired by the abstractionist approach outlined in Section 3. Recall that this approach is concerned with *abstraction terms*, which were defined as singular terms of the form '$\Sigma(\alpha)$', where $\Sigma$ is an operator associated with an equivalence relation $\sim$ and $\alpha$ is some item in the domain of this equivalence relation. For instance, a neo-Fregean number term is of the form '$\#F$', where $\#$ is the cardinality operator and $F$ is a concept. So here the abstractionist approach be reflected in the syntactic structure of the singular terms in question.

I wish to avoid this problematic claim about syntactic structure. My version of abstractionism will instead be a matter of how singular terms with no internal syntactic structure come to possess semantic values. So instead of abstraction terms I will be concerned with syntactically simple singular terms which are associated with an item $\alpha$ and a relation $\sim$ defined on such items. For instance, in the language of arithmetic, each singular term is asso-

[31] See e.g. (Stalnaker, 1997). My distinction between semantics and meta-semantics is thus the same as Stalnaker's distinction between "descriptive" and "foundational" semantics. See e.g. (Stalnaker, 2001).

ciated with a concept $F$ and the equivalence relation of equinumerosity; and in the language of directions, each singular term is associated with a line $l$ and the equivalence relation of parallelism. The item associated with a singular term must not be confused with the term's referent. The role of these items is rather to present the referents. For instance, in the case arithmetic, a concept serves to present the cardinal number of the objects falling under the concept; and in the case of directions, a line serves to present the direction that it has. Let's refer to the items that play this role of presenting the proper referents as *presentations*. I will mostly use lower-case Greek letters to range over presentations.

The role of the relation $\sim$ is to specify when two presentations determine the same referent. Let's refer to relations that play this role as *unity relations*. This means that a unity relation has to be symmetric: for if two presentations $\alpha$ and $\beta$ determine the same referent, then so do $\beta$ and $\alpha$. It also means that a unity relation has to be transitive: for if $\alpha$ and $\beta$ determine the same referent and $\beta$ and $\gamma$ do so as well, then clearly so must $\alpha$ and $\gamma$. A unity relation is typically also assumed to be reflexive. This assumption will later be challenged but will be left in place for now. For the time being I will thus be assuming that unity relations are equivalence relations.

Recall that the abstractionist view claims that it suffices for an abstraction term to refer that it has sense and that its presentation stands in the relevant unity relation to itself. What does this claim correspond to in my slightly modified approach? Presumably a singular term is guaranteed to have sense already by the fact that it is associated with a presentation and a unity relation. The abstractionist view can then be characterized simply as the view that it suffices for a term to refer that it has been assigned a presentation and a unity relation. In what follows I develop and defend a refinement of this view. I first develop two examples.[32] Then I state the sufficient condition for possession of semantic value in greater generality.

### 7.2 The case of directions

Let $D$ be a domain of lines and other directed items. Assume $\mathcal{L}$ is a first-order language with identity such that:

(i) the variables of $\mathcal{L}$ range over $D$,

[32] The examples are only meant to illustrate the sufficient condition. One or both of the examples can thus be rejected without thereby rejecting the sufficient condition.

(ii) each singular term of $\mathcal{L}$ has been assigned an element of $D$,

(iii) each atomic predicate of $\mathcal{L}$ is defined on each element of $D$,

(iv) for any two singular terms $t_1$ and $t_2$ that have been assigned $l_1$ and $l_2$ we have:

$$\ulcorner t_1 = t_2 \urcorner \text{ is true} \quad \text{iff} \quad l_1 \parallel l_2.$$

When these assumptions are met, I say that $\mathcal{L}$ has a *pre-interpretation.* (A more formal definition will be provided later.) A pre-interpretation is much like a proper interpretation except that it is based on a domain of presentations (in this case lines) instead of proper referents (in this case directions). A consequence of this is that the identity predicate is interpreted non-standardly: the identity predicate can be true of two non-identical lines provided they are parallel.

Note that the standard laws of identity require that every predicate $\mathbf{P}$ of $\mathcal{L}$ be a congruence with respect to parallelism, in the sense that the following holds:

> Assume $\mathbf{P}$ is $n$-adic and that $l_i \parallel l'_i$ for each $i$ from 1 to $n$. Then $\mathbf{P}$ holds of $l_1, \ldots, l_n$ iff $\mathbf{P}$ holds of $l'_1, \ldots l'_n$.

My sufficient condition says roughly that a pre-interpretation suffices for a proper interpretation. But let's be more precise. Assume $\mathcal{L}$ has a pre-interpretation and that $t_i$ are singular terms of $\mathcal{L}$ which have been assigned lines $l_i$ respectively. Then the sufficient condition says that expressions from $\mathcal{L}$ can be assigned semantic values such that:

(a) singular terms have the same semantic value iff their assigned lines are parallel, that is:

$$[\![t_1]\!] = [\![t_2]\!] \leftrightarrow l_1 \parallel l_2$$

(b) the principle of compositionality holds for simple predications, that is:

$$[\![\mathbf{P}(t_1, \ldots, t_n)]\!] = [\![\mathbf{P}]\!]([\![t_1]\!], \ldots, [\![t_n]\!])$$

(c) the semantic values are *sui generis.*

Some explanations are in order. The first claim, (a), says that the semantic value assigned to a singular term $t_i$ depends only on the assigned line $l_i$ up to parallelism. There is thus a function $d$ such that the semantic values of the terms $t_i$ are given as

$$(3) \qquad [\![t_i]\!] = d(l_i)$$

and such that

$$(4) \qquad d(l_i) = d(l_j) \leftrightarrow l_i \parallel l_j.$$

Note that this latter formula is just an abstraction principle. So we now have an explanation of why such principles are so important: they play an important role in the meta-semantic account of reference.

To understand the second claim, (b), let's consider an example. Assume $\mathcal{L}$ contains a two-place predicate **P** which is true of $t_i$ and $t_j$ iff $l_i$ and $l_j$ are orthogonal. Then (b) says that **P** has a semantic value $[\![\mathbf{P}]\!]$ which is true of the directions $d(l_i)$ and $d(l_j)$ iff the lines $l_i$ and $l_j$ are orthogonal. This means that $[\![\mathbf{P}]\!]$ is an orthogonality predicate on directions. We are here relying on the fact that two directions are orthogonal iff any two lines whose directions they are, are orthogonal.

Claims (a) and (b) are highly plausible. To see this, recall what semantic values are supposed to do. The semantic value of an expression was explained as the contribution that the expression makes to the truth-value of sentences in which it occurs. Assume that $\mathcal{L}$ has a pre-interpretation and that the singular term $t$ has been assigned a line $l$ as its presentation. What is the semantic contribution that $t$ makes to the truth-values of sentences in which it occurs? Clearly $t$ makes *some* contribution: for all sentences involving the term have truth-values which typically depend on the presentation $l$. But the contribution cannot be anything as specific as the line $l$. For we know that any parallel line $l'$ would make precisely the same contribution. Rather, the semantic contribution of $t$ must be something which is shared by all lines $l'$ that are parallel with $l$. But this is exactly the sort of contribution that claim (a) ascribes to $t$. An analogous motivation can be provided for (b). Clearly the atomic predicates make some semantic contribution. But this contribution does not discriminate between singular terms with the same semantic value, as asserted by (b).

In pure mathematics it would be natural to represent the semantic contribution of the singular term $t$ as the equivalence class of $l$ under the equivalence relation of parallelism, that is, as the set of all lines $l'$ that are parallel to $l$. This too would ensure that claim (a) holds. Moreover, we could let the semantic value of the predicate for orthogonality be the function that maps two such equivalence classes to the true iff any two lines from each of the two classes are orthogonal, and to the false otherwise. This is easily seen to ensure that claim (b) holds. So this provides a useful model of the desired assignment of semantic values.

However, a model of an assignment is not the same as the intended assignment itself. Although the semantic values of the language of directions can be *represented by* equivalence classes of lines, they should not be *identified with* such equivalence classes. Doing so would ascribe to the semantic values properties that go beyond the contribution that the relevant terms make to the truth-values of sentences in which they occur. For instance, equivalence classes have set theoretic elements, which is a notion that is completely foreign to the geometrical language in question. Accordingly the third claim of the sufficient condition, (c), says that it is permissible to assign to the expressions in question semantic values that are primitive and *sui generis* and not just set theoretic constructs. These semantic values are nothing more than the contributions made by the relevant singular terms.

### 7.3 The case of physical bodies

The second example is reference to physical bodies. How does such reference come about? Since human beings are very complex organisms who stand in very complex social relations, it will be useful to simplify. What we want is a good model of this form of reference, not necessarily a description that is accurate in every detail. I will therefore attempt to develop a model based on robots or computers embedded in, and interacting with, a physical environment. What is required of a robot for it to refer to physical bodies, such as sticks and stones, in its environment? We get a better understanding of the problem by focusing on the senses of sight and touch, and on some very fundamental thought processes. Other senses, such as smell, taste, and hearing, appear to play a less fundamental role in our reference to physical bodies. Consciousness too (in the sense of awareness of what it is like to have various sorts of experiences) will be put to one side, as it appears inessential to our core notion of reference.

So consider a robot equipped with senses of sight and touch similar to our natural human senses. Such a robot must interact with its environment by detecting light reflected by surrounding surfaces and by having a capacity for touching and grasping things in its vicinity. What is required for such a robot to make reference to a physical body in its environment? Obviously the robot must "perceive" the body in the sense that it must receive light from some part of its surface or touch some part of it. The robot will thus receive information from some spatiotemporal part of the body. These parts need not have natural boundaries in either space or time; they are simply the sum-totals of the particle-instants with which the robot causally interacts in this perception-like way.

But it is not sufficient for a robot to refer to a body that it should receive such information from some part of the body. The robot also needs a mechanism for determining when two such pieces of information belong to the same body. This task is far from trivial. For we are surrounded by bodies that are partially hidden, that are occluded by other bodies, and that move in and out of view. For instance, a stick can be partially buried, and a stone can be partially covered by other stones piled up around it. So there will always be different ways of "getting at" one and the same physical body, both from different spatial points of view and at different moments of time. It is therefore essential that the robot have some mechanism for grouping together pieces of perceptual information that belong to one and the same body.

I claim that what matters for this task is that the parts from which the robot receives perceptual information be spatiotemporally connected in some suitable way. Assume for instance that the robot establishes visual contact with part of a stick that emerges from the ground and that one of its "arms" is simultaneously probing into the ground nearby and encountering something hard. What should we "teach" the robot about the conditions under which the two parts it interacts with belong to the same body? Roughly, the kind of connectedness that matters has to do with solidity and motion: The two parts must be related through a continuous stretch of solid[33] stuff, all of which belongs to the same unit of independent motion (roughly in the sense that, if you wiggle one part, the other part follows along). A more precise answer will ultimately be needed. But this will have to suffice for present purposes.

I believe this relatively simple model captures the core of the phenomenon of reference to physical bodies. What matters is that our agents (whether human or robot) receive sensory

[33]I here mean 'solid' in the ordinary sense in which a stick or a stone is said to be solid. Of course, physics tells us that even sticks and stones aren't solid in the stricter sense of filling up all space at an atomic level.

information from parts of bodies and that they have a capacity for grouping together such pieces of information just in case these pieces derive from parts that are spatiotemporally connected in the way spelled out above. The parts in question serve as presentations of bodies. An agent's relation to these presentations is typically purely causal. Next, let $\sim$ be this relation of spatiotemporal connectedness. This serves as a unity relation on the presentations.

Let $\mathcal{L}$ be a first-order language with identity. Assume the singular terms and variables of $\mathcal{L}$ refer to and range over the above sort of presentations. Assume further that all predicates of $\mathcal{L}$ are defined on these presentations and that an identity statement $\ulcorner t_1 = t_2 \urcorner$ comes out true just in case the assigned presentations stand in the relevant unity relation. Then I claim that there is a partial function $B$ that maps a part $\alpha$ to the physical body, if any, that $\alpha$ picks out. That is, $B(\alpha)$ is the body that $\alpha$ is part of. Bodies are then subject to the following "abstraction principle":

$$B(\alpha) = B(\beta) \leftrightarrow \alpha \sim \beta \tag{5}$$

The case of physical bodies is thus analogous to that of directions; in particular, a pre-interpretation suffices for a real interpretation. The most important difference is that in the case of bodies some presentations may fail to determine referents. For instance, if I point to ground, I may be causally related with a presentation which fails to determine a unique physical body. This requires a slight modification of our approach. Recall that a unity relation is supposed to tell when two presentations determine the same referent. We used this to motivate the symmetry and transitivity of unity relations. However, no motivation was provided for the commonly made assumption that unity relations are reflexive as well (and thus are equivalence relations). We now see that the assumption of reflexivity is problematic. For if a presentation $\alpha$ fails to determine a referent, it is false that $\alpha$ determines the same referent as $\alpha$. So we cannot assume that $\sim$ is an equivalence relation. All we know is that $\sim$ is an equivalence relation on its field, that is on the class of items $\alpha$ which are related by $\sim$ to some $\beta$. Accordingly we must also allow the function that maps presentations to their referents to be partial (as in (5) above).

Let's call an ordered pair of a presentation $\alpha$ and a unity relation $\sim$ a *referential attempt*. And let's say that a referential attempt $\langle \alpha, \sim \rangle$ *is successful* iff $\alpha \sim \alpha$.

### 7.4 A general form of the sufficient condition

I now formulate a more general form of the sufficient condition for possession of semantic value.

**Definition 1** Say that a first-order language $\mathcal{L}$ has a **pre-interpretation** iff there is a domain $D$ of successful referential attempts, based on one and the same unity relation $\sim$, such that:

(i) the variables of $\mathcal{L}$ range over $D$,

(ii) each singular term of $\mathcal{L}$ has been assigned a referential attempt from $D$,

(iii) each atomic predicate of $\mathcal{L}$ is defined on all referential attempt from $D$,

(iv) for any two singular terms $t_1$ and $t_2$ that have been assigned presentations $\alpha_1$ and $\alpha_2$ respectively, we have:

$$\ulcorner t_1 = t_2 \urcorner \text{ is true} \quad \text{iff} \quad \alpha_1 \sim \alpha_2.$$

As in the case of directions, it follows from the laws of identity that every predicate of $\mathcal{L}$ is a congruence with respect to the relation $\sim$.

**Sufficient Condition for Possession of Semantic Values**. Let $\mathcal{L}$ be a first-order language with identity. Assume $\mathcal{L}$ has a pre-interpretation. Then expressions of $\mathcal{L}$ can be assigned semantic values such that:

(a) two singular terms $t_1$ and $t_2$ have the same semantic value iff their associated referential attempts $\langle \alpha_1, \sim \rangle$ and $\langle \alpha_2, \sim \rangle$ are equivalent in the sense that $\alpha_1 \sim \alpha_2$; that is:

$$[\![t_i]\!] = [\![t_j]\!] \leftrightarrow \alpha_i \sim \alpha_j$$

(b) the principle of compositionality holds for simple predications; that is:

$$[\![\mathbf{P}(t_1, \ldots, t_n)]\!] = [\![\mathbf{P}]\!]([\![t_1]\!], \ldots, [\![t_n]\!])$$

(c) the semantic values are *sui generis*.

In short, the sufficient condition says that a pre-interpretation suffices for the existence of a proper interpretation.

As before, (a) ensures that the semantic value assigned to a singular term $t_i$ can be determined as a function of the presentation $\alpha_i$, where this function is invariant under the unity relation $\sim$. That is, there is a function $f$ such that

(6) $$[\![t_i]\!] = f(\alpha_i)$$

and such that the following abstraction principle holds:

(7) $$f(\alpha_i) = f(\alpha_j) \leftrightarrow \alpha_i \sim \alpha_j.$$

An initial defense of the sufficient condition has been developed in previous subsections. This defense is based on the observation that when a language has a pre-interpretation, it behaves precisely *as if* its expressions have semantic values in accordance with these two claims. But this makes it hard to deny that the expressions really have semantic values of this sort. This defense involves a kind of reversal of the explanatory direction usually associated with the principle of compositionality. Usually the assignment of semantic values to simple expressions is taken for granted, and the principle of compositionality is used to determine and explain the assignment of semantic values to complex expressions. But here we have proceeded in the opposite direction. We began by assuming that the language in question has a pre-interpretation. This means that the truth-values of simple predications are determined by the presentations assigned to the singular terms. But it also means that these truth-values depend on the presentations only up to the equivalence given by the unity relation. On this basis I claimed that the expressions of the language in question have semantic values in accordance with the sufficient condition. For what it *is* for an expression to have a semantic value is just for the expression to make this kind of systematic contribution to the truth-values of sentences in which it occurs.

However, the defense of the sufficient condition is not yet complete. For what are we to say to a skeptic who grants that a language has a pre-interpretation but who doubts that this suffices for the existence of a real interpretation? The skeptic may refuse to accept the existence of the semantic values on which the interpretation is based. This takes us back to the Problem of Existence, described in Section 4.

## 8 The Problem of Existence

My strategy for addressing this problem will be to turn the *sui generis* character of the desired semantic values to my own advantage. This *sui generis* character means that any characterization of the semantic values must ultimately be based on the relevant sort of referential attempts. And this in turn means that any characterization in the meta-language of the assignment of semantic values to expressions of the object language will have to make use of terms whose reference is ultimately based on the relevant sort of referential attempts. This will have important consequences for the ensuing discussion. Let $t$ be a singular term of the object language which is associated with a successful referential attempt $\langle \alpha, \sim \rangle$. Consider now the question (formulated in the meta-language) whether $t$ has a unique semantic value. There are two kinds of case to consider. In *the simple case* we assume that whenever a singular term of the meta-language is associated with a successful referential attempt, it refers. In *the hard case* this assumption is not granted.

In the simple case our question allows of a straightforward answer. Let's begin with *the existence* of a semantic value of $t$. To establish this it suffices to find a meta-language term $a$ which corefers with $t$. But in the simple case this is easy: we just consider some meta-language term $a$ which is associated with a referential attempt $\langle \beta, \sim \rangle$ such that $\alpha \sim \beta$. Our account then entails that $a$ co-refers with $t$.

*The uniqueness* of $t$'s semantic value is similarly easy to establish. Assume some philosopher studies the object language and asks which of two candidate referents the term $t$ denotes. Since the relevant semantic values are assumed to be *sui generis*, these two candidate referents must themselves be picked out by means of the relevant sort of referential attempts. In particular, since two referential attempts can pick out the same referent only if their unity relations are the same, these referential attempts must be of the form $\langle \beta_i, \sim \rangle$. But then our philosopher can determine which (if any) candidate referent is the right one simply by determining for which (if any) $i$ we have $\alpha \sim \beta_i$.

It will be objected that this defense of the (unique) existence of a semantic value simply begs the question. Any doubt about the existence of a semantic value for the object language term $t$ will translate into an analogous doubt about the existence of a semantic value for the meta-language term $a$. When a skeptic doubts that it suffices for the object language term $t$ to refer that it is associated with a successful referential attempt, this doubt will apply with equal force to the analogous assumption concerning the meta-language term $a$. In short, the skeptic will object that we are not entitled to assume that the meta-language is capable of the sort of reference whose existence in the object language is being disputed.

This objection forces us to consider the hard case, where the mentioned assumption concerning the meta-language is withheld. Then we have no guarantee that the desired assignment of semantic values can be described in the meta-language. For since the intended semantic values are *sui generis*, this assignment can only be described using referential resources of the meta-language which in the hard case are no longer assumed to be available. How damaging is this? A comparison may be useful. Assume the object language has a pre-interpretation of the sort associated with talk about physical bodies, as described in the previous section. Based on this we want to claim that the terms of the language succeed in referring to physical bodies. However, no such claim can be made if our meta-language lacks the resources to refer to physical bodies. We cannot adequately describe the semantics of a language that purports to talk about physical bodies without ourselves referring to physical bodies. But it would be crazy to take this as an indication that the object language fails to refer to physical bodies. The problem should rather be blamed on the impoverished nature of the meta-language.

I suggest that a similar response is appropriate in our hard case. It is trivial that in a meta-language that does not allow for reference to some sort of objects, say $F$s, one cannot adequately describe an interpretation which assigns to some object language $F$s as semantic values. But the appropriate response to this trivial observation is not to conclude that the object language fails to refer to $F$s but rather that the meta-language will have to be expanded if it is to be capable of adequately describing the semantics of this object language. That is, we must allow terms of the meta-language to be associated with the successful referential attempts of the sort in question and grant that this suffices for the meta-language terms to refer. But once such an expansion is carried out, we are back in the simple case.

The crucial question is thus whether this sort of expansion is legitimate. I claim that it is, whereas the skeptic will challenge this. *Prima facie* my claim appears more plausible. For the expansion I have described is a natural, systematic, and provably consistent way of using language. Moreover, the expansion makes available a description of agents and their language in terms of semantic values and compositionality, without having to endow these semantic values with any excess structure, as would otherwise be the case (for instance if they were taken to be set-theoretic equivalence classes). This sort of semantic description is something we engage in all the time: in daily life as well as in more scientific accounts of human action, and applied to discourses about physical objects as well as abstract ones. In short, in order to give the best explanation of the semantics the object language, the meta-language must be capable of precisely the kind of reference in question.

What arguments can the skeptic give against the legitimacy of the described expansion which would overrule this *prima facie* defense of it? One option is to argue that the expansion remains impermissible until it has been established that its new referential resources really succeed in referring. But this demand would be unreasonable. The demand could only be satisfied in some meta-meta-language. But how do we know that the terms of this language refer? In order to ensure that they do, we would have to ascend to yet another language, thus embarking on an infinite regress. In fact, since the demand is completely general and not limited to the disputed kind of objects, this regress would arise for all kinds of objects, including those that the skeptic accepts.

A better option for the skeptic is to protest on the ground that no referent is present in the proposed basis for the ascription of a semantic value to an object language term $t$. Assume that $t$ is associated with the referential attempt $\langle \alpha, \sim \rangle$. Then the proposal is that it suffices to find a meta-language term $a$ associated with a referential attempt $\langle \beta, \sim \rangle$ such that $\alpha \sim \beta$. But no referent is mentioned in this proposed basis. So how could this proposed basis possibly be sufficient for the existence of a referent?

In some cases the answer is that the referent plays an *implicit* role in the fact that $\alpha \sim \beta$. Consider the case of physical bodies. Let $\alpha$ and $\beta$ be two chunks of spatiotemporal stuff. What is required for $\alpha \sim \beta$ is then that the two chunks belong to a totality of physical stuff with sufficiently sharp and natural boundaries and which moves as a unit. So in this case there is nothing odd about the claim that $\alpha \sim \beta$ ensures the existence of a physical body, which serves as the referent of the term $t$.

Things are less clear in cases involving mathematical objects. Consider the case of directions. I argued that directions are presented by means of lines and subject to parallelism as the unity relation. Assume that two lines $l_1$ and $l_2$ are parallel. Does it make sense to say that a referent—a direction—plays an implicit role in the fact that $l_1$ is parallel to $l_2$? Given how little is involved in this fact, it is far from obvious that it does. In fact, since every line is parallel to itself, every lines gives rise to a successful referential attempt. Can success really be that cheap?

The correct response is, I believe, not to conclude that a language with a pre-interpretation may nevertheless fail to possess appropriate semantic values, but rather to admit that semantic values can be extremely thin. In the case of directions a reductionist interpretation of the semantic values may even be appropriate. If the existence of a direction requires nothing beyond the trivial fact that the associated line is parallel to itself, then this may be a form of reductionism. But if so, the correct response is not to conclude that there are no directions but rather that the existence of directions reduces to facts about lines. By comparison, assume that the mind will one day be found to reduce completely to the physical. The correct response to this discovery would not be to conclude that people have no minds but rather that the existence of minds reduces to facts about the physical.

My aim has been to explain how singular terms can come to possess semantic values, not to defend some robust non-reductionist account of these semantic values. What is required for the existence of a semantic value will vary enormously from one type of entity to another. I have been gesturing at this variation throughout the article in my informal and metaphorical talk of objects as thick and thin. I have suggested that physical bodies are thick because their existence makes substantive demands on the world, whereas natural numbers are thin because their existence is very cheap. In the next section I suggest two ways of giving more precise content to the intuitive contrast between thick and thin objects.

## 9 The varieties of semantic value

The most obvious respect in which an object can be thick is by having spatiotemporal location. This typically also ensures causal efficacy. Such objects are known as *concrete*. At the other extreme are *pure abstract* objects, such as the natural numbers, which are non-spatiotemporal and without any kind of intrinsic relation to space or time. Concrete objects are thicker than pure abstract ones because their existence makes substantive demands on the particular region of spacetime where they are located. By contrast, the existence of a pure abstract object makes no demands on spacetime at all. In between the concrete and pure abstract objects are what Charles Parsons calls *quasi-concrete* objects, that is, objects which despite being abstract have canonical realizations in spacetime.[34] For instance, directions and geometrical figures have canonical realizations in the form of tokens and concrete figures with the shapes in question. Quasi-concrete objects are thus somewhat thicker than pure abstract objects in that their existence makes some demands upon spacetime, however general and indirect. There can for instance be no quasi-concrete twenty-dimensional geometrical figures because spacetime isn't rich enough to allow canonical realizations of such figures.

Another respect in which an object can be thick is by having properties which go beyond those that are implicitly contained in any one of its presentations. Consider for instance the question whether a physical body has the property of being a solid sphere. This question cannot be answered on the basis of any one of the body's proper parts; rather, information will be needed about the entire body. And there is nothing unusual about this case. Whether a body has some property will in general depend on all or many of its parts. A body therefore has many of its properties in an irreducible way, that is, in a way that isn't reflected in any property of any one of its proper parts. And this means that physical bodies play an ineliminable role in making predications true: for the truth of such predications cannot be reduced to a matter of how the body is presented.

The natural numbers are very different in this regard. Consider the question whether a natural number $n$ has the property of being even. Unlike the case of a physical body's being a solid sphere, any concept $F$ by means of which the number $n$ is presented suffices to answer the question:[35] just examine whether $F$ has an even number of instances. There is no need to examine other presentations of $n$ (or the number itself, whatever exactly that would involve). I argue elsewhere that this observation generalizes: for any arithmetical property $A$, the question whether $n$ possesses $A$ can be reduced to a question about any concept by means of which $n$ is presented.[36] A natural number is in this way "impoverished" compared to the concepts that present it, as all of its properties are already implicitly contained in each of these concepts. This provides further support for the reductionist account of the natural numbers suggested at the end of the previous section. When an object is thin in this respect, it becomes hard to resist the idea that they are mere "shadows of syntax" (in the apt metaphor of (Wright, 1992), pp. 181-2). So my view here contrasts starkly with that of

[34] See (Parsons, 1980).
[35] Recall that a natural number $n$ is said to be presented by a concept $F$ iff $n = \#F$.
[36] See [???].

the neo-Fregeans, who insist that abstraction principles give access to objects "every bit as objective as mountains, rivers and trees" ((Wright, 1983), p. 13).

There may be a variety of reasons to be interested in objects that are thick in either of the two respects just described or in others. But it would be a mistake to regard thickness in any of these respects as a necessary condition for objecthood in general. In particular, doing so would beg the question against the meta-ontological minimalist, whose claim is only the modest one that certain sentences have semantic values, not that the objects that serve as semantic values are in any way thick.

A more interesting question is whether some requirement of thickness is needed to support a platonistic understanding of some class of abstract objects. My view has some of the key features of platonism. For instance, on my view it is true to say that there exist abstract mathematical objects. And this claim really has the semantic structure that it appears to have, namely the one it shares with less problematic existence claims. Nevertheless it is natural to think that the semantic values that make the above ontological claim true are so lightweight as to make the label 'platonism' inappropriate.

In fact, platonism is often defined as the conjunction of the already mentioned view that there are abstract objects and some claim to the effect that these objects are "fully real" or "fully independent" in their existence. It is far from obvious how this second conjunct should be articulated. But if some plausible articulation could be provided and shown to go beyond what my minimalism delivers, I would happily admit that the relevant objects should not to be understood along traditional platonist lines. This admission would be fully compatible with the goal of this article, which is only to defend the first conjunct of the definition of platonism, not the second. And it is this first conjunct that holds the key to the benefits of minimalism described in Section 2.

Recall the four problems described in Section 4 which face any form of minimalism. The Problem of Existence has been addressed above and in the previous two sections. But our discussion has also developed the resources that are needed to address another two of the problems.

First there is the Problem of Overgeneration, which asks how one can accept the idea of thin objects without being led to the absurd conclusion that all objects are thin. For instance, why is so much more required for the existence of physical bodies than for the existence of natural numbers? The answer emerging from the above discussion is clear: some objects are naturally thicker than others. How thick some kind of object is depends on the nature of the presentations and the unity relation with which this kind is associated. For instance, for there to be a physical body on my desk, there must be some chunk of matter on the desk which is spatiotemporally connected in a suitable natural and continuous way to other chunks of matter. This is quite a demanding condition. Other kinds of objects are naturally much thinner. For instance, the existence of a direction requires nothing more than the (possible) existence of an appropriately oriented line.

Next there is the Problem of Uniformity, which asks whether there is enough uniformity between thick and thin objects for both to be classified as objects. Again an answer emerges from the above discussion. The common structure is that provided by the notion of semantic value. Despite their dramatic differences, both thick and thin objects serve as semantic values of singular terms. And if we accept the Fregean slogan that to be is to be a possible referent of a singular term, then both thick and thin semantic values are rightly regarded as objects.

What about Dummett's view that reference to thin objects is "semantically idle" and thus not a genuine form of reference? On my account singular terms have semantic values just as much in cases of thin objects as in cases of thick ones. And these semantic values play the same kind of role in a compositional semantics in both cases. I am therefore opposed to any reductionist or deflationist understanding of the *semantics* of sentences concerned with thin objects, which is what Dummett appears to defend. To the extent that my view allows any form of reductionism, this will be at the *meta-semantic* level. For I am willing to admit that the facts which make it true that a singular term refers to some thin object may not themselves involve this object.

Finally there is the Problem of Consistency, which asks whether minimalism can avoid inconsistency. In fact the version of minimalism developed here seems particularly vulnerable to inconsistency. I have argued that when a first-order language has a pre-interpretation, it also has a real interpretation. I am therefore committed to the view that abstraction is permitted on *any* equivalence relation on *any* domain of entities. But consider the domain consisting of all concepts. And consider some equivalence relation on this domain which is sufficiently fine-grained, for instance that of co-extensionality. Then cardinality considerations appear to show that there cannot be enough objects to serve as abstracts for all of the resulting equivalence classes.

Although it would take me too far afield to attempt to solve this problem here, I owe the reader at least a sketch of what a solution might look like.[37] The key is to observe that the domains on which abstraction is dangerous are those of concepts and other higher-order entities. But although I am committed to abstraction on any equivalence relation on any domain of entities, I have not made any substantive commitments concerning what concepts and other higher-order entities there are. The threat of inconsistency arises only when one takes a liberal view of the matter. But my minimalism supports a more restrictive view. Just as an object is a possible semantic value of a singular term, a concept is a possible semantic value of an open formula. The question thus arises what is required for an open formula to possess a semantic value. I defend a fairly restrictive answer to this question and prove that the resulting view of concepts and abstraction is consistent.[38]

Can the other forms of minimalism described in Section 3 produce equally good answers to the four problems facing any minimalist view? It is far from clear that they can. The coherentist and priority-based forms of minimalism lack some of the key resources employed in my defense of abstractionist minimalism. In particular, they lack the resources to provide an informative account of how singular terms and first-order variables come to refer to and range over appropriate semantic values.

## References

Balaguer, M. (1998). *Platonism and Anti-Platonism in Mathematics.* Oxford University Press, Oxford.

Beaney, M. (1997). *The Frege Reader.* Blackwell, Oxford.

Benacerraf, P. (1973). Mathematical Truth. *Journal of Philosophy*, 70(19):661–679.

Boolos, G. (1990). The Standard of Equality of Numbers. In Boolos, G., editor, *Meaning and Method: Essays in Honor of Hilary Putnam.* Harvard University Press, Cambridge, MA. Reprinted in (Boolos, 1998).

Boolos, G. (1997). Is Hume's Principle Analytic? In Heck, R., editor, *Logic, Language, and Thought.* Oxford University Press, Oxford. Reprinted in (Boolos, 1998).

[37] See (Linnebo, a) for a beginning of an answer.

[38] I regard it as an advantage of this response to the Problem of Inconsistency that it makes crucial use of my version of minimalism. By contrast, Hale and Wright regard the problem as independent of their account of what abstraction principles do.

Boolos, G. (1998). *Logic, Logic, and Logic.* Harvard University Press, Cambridge, MA.

Dummett, M. (1956). Nominalism. *Philosophical Review*, LXV:491–505. repr. in (Dummett, 1978).

Dummett, M. (1978). *Truth and Other Enigmas.* Harvard University Press, Cambridge, MA.

Dummett, M. (1981). *Frege: Philosophy of Language.* Harvard University Press, Cambridge, MA, second edition.

Dummett, M. (1991). *Frege: Philosophy of Mathematics.* Harvard University Press, Cambridge, MA.

Eklund, M. Neo-Fregeanism and Metaontology,.

Eklund, M. (2006). Neo-Fregean Ontology. In Hawthorne, J., editor, *Philosophical Perspectives*, volume 20: Metaphysics. Blackwell.

Ewald, W. (1996). *From Kant to Hilbert: A Source Book in the Foundations of Mathematics*, volume 2. Oxford University Press, Oxford.

Field, H. (1984). Platonism for Cheap? Crispin Wright on Frege's Context Principle. *Canadian Journal of Philosophy*, 14:637–62. Reprinted in (Field, 1989).

Field, H. (1989). *Realism, Mathematics, and Modality.* Blackwell, Oxford.

Frege, G. (1891). Function and concept. Reprinted in (Beaney, 1997).

Frege, G. (1892). Concept and object. Reprinted in (Beaney, 1997).

Frege, G. (1953). *Foundations of Arithmetic.* Blackwell, Oxford. Transl. by J.L. Austin.

Frege, G. (1980). *Philosophical and Mathematical Correspondence.* University of Chicago Press, Chicago. ed. by G. Gabriel et al., transl. by H. Kaal.

Hale, B. and Wright, C. The *Metaontology* of Abstraction. In et al., D. C., editor, *Metametaphysics.* Oxford University Press, Oxford.

Hale, B. and Wright, C. (2001). *Reason's Proper Study.* Clarendon, Oxford.

Hale, B. and Wright, C. (2008). Focus Restored: Comments on John MacFarlane. *Synthese*. Forthcoming.

Hellman, G. (2001). Three Varietites of Mathematical Structuralism. *Philosophia Mathematica*, 9(3):184–211.

Kant, I. (1997). *Critique of Pure Reason*. Cambridge University Press. Transl. by Paul Guyer and Allen Wood.

Keränen, J. (2001). The Identity Problem for Realist Structuralism. *Philosophia Mathematica*, 9(3):308–330.

Lewis, D. (1991). *Parts of Classes*. Blackwell, Oxford.

Linnebo, Ø. Bad Company Tamed. forthcoming in *Synthese*.

Linnebo, Ø. Introduction. forthcoming in *Synthese*.

Linnebo, Ø. (2006). Epistemological Challenges to Mathematical Platonism. *Philosophical Studies*, 129(3):545–574.

Parsons, C. (1965). Frege's Theory of Number. In Black, M., editor, *Philosophy in America*. Cornell University Press. Reprinted in (Parsons, 1983).

Parsons, C. (1980). Mathematical Intuition. *Proceedings of the Aristotelian Society*, 80:145–68.

Parsons, C. (1983). *Mathematics in Philosophy*. Cornell University Press, Ithaca, NY.

Parsons, C. (1990). The Structuralist View of Mathematical Objects. *Synthese*, 84:303–346.

Quine, W. (1986). *Philosophy of Logic*. Harvard University Press, Cambridge, MA, second edition.

Rayo, A. and Uzquiano, G., editors (2006). *Absolute Generality*, Oxford. Oxford University Press.

Resnik, M. (1997). *Mathematics as a Science of Patterns*. Oxford University Press, Oxford.

Rosen, G. and Dorr, C. (2002). Composition as fiction. In Gale, R., editor, *Blackwell Guide to Metaphysics*, pages 151–174. Blackwell, Malden.

Shapiro, S. (1997). *Philosophy of Mathematics: Structure and Ontology.* Oxford University Press, Oxford.

Sider, T. (2007). Parthood. *Philosophical Review*, 116:51–91.

Stalnaker, R. (1997). Reference and Necessity. In Hale, B. and Wright, C., editors, *Blackwell Companion to the Philosophy of Language*, pages 534–554. Blackwell, Oxford.

Stalnaker, R. (2001). On Considering a Possible World as Actual. *Proceedings of the Aristotelian Society*, Suppl. vol. 65:141–156.

Wright, C. (1983). *Frege's Conception of Numbers as Objects.* Aberdeen University Press, Aberdeen.

Wright, C. (1992). *Truth and Objectivity.* Harvard University Press, Cambridge, MA.