

To appear in *Philosophy and Phenomenological Research*

## QUANTIFICATION AND REALISM\*

Michael Glanzberg

Department of Linguistics and Philosophy  
MIT  
Cambridge, MA 02139  
glanzber@mit.edu

March, 2003

---

\*Thanks to Richard Cartwright, Warren Goldfarb, Richard Heck, Vann McGee, Charles Parsons, Agustín Rayo, Ralph Wedgwood, Timothy Williamson, and an anonymous referee for helpful comments and discussions. Versions of this paper were presented at the University of California, Santa Cruz, MIT, and McGill University. I am grateful to my audiences there for valuable discussions.

I owe a special debt of gratitude to the late George Boolos. This paper grew out of one I started long ago, while I was a graduate student. George kindly encouraged the project, and was wonderfully generous with his time and his insight. All this in spite of his disagreeing with practically everything I said. It was a privilege, and a pleasure, to have learned from him.

## **Abstract**

This paper argues for the thesis that, roughly put, it is impossible to talk about absolutely everything. To put the thesis more precisely, there is a particular sense in which, as a matter of semantics, quantifiers always range over domains that are in principle extensible, and so cannot count as really being ‘absolutely everything’. The paper presents an argument for this thesis, and considers some important objections to the argument and to the formulation of the thesis. The paper also offers an assessment of just how implausible the thesis really is. It argues that the intuitions against the thesis come down to a few special cases, which can be given special treatment. Finally, the paper considers some metaphysical ideas that might surround the thesis. Particularly, it might be maintained that an important variety of realism is incompatible with the thesis. The paper argues that this is not the case.

In this paper, I shall attempt to accomplish two things. One is to argue for a thesis that many find implausible. To put it roughly, and bit melodramatically, it is impossible to talk about absolutely everything. To put it more precisely, there is a particular sense in which, as a matter of semantics, quantifiers always range over domains that are in principle extensible, and so cannot count as really being ‘absolutely everything’. The second thing I shall try to accomplish is to assess just how implausible my position really is. I shall suggest that the intuitions against it come down to a few special cases, which can be given special treatment. I shall also consider some metaphysical ideas that might surround the view. Particularly, it might be maintained that an important variety of *realism* is incompatible with my position. I shall argue this is not the case, and examine what the connections between quantification and realism might then be.

Before starting out, a couple of preliminaries need to be addressed. First, I have yet to give an accurate formulation of my principle claim. Instead, I put it in a rather cagey way, saying, ‘in a sense’ and ‘as a matter of semantics’. However, the background needed to make these provisos precise is part of the argument I shall present, so I shall delay giving the proper formulation of my thesis until they have been explained. I must also beg an indulgence.

Following the patron bishop of those who would argue for implausible views, I shall often speak with the vulgar, and talk about ‘absolutely everything’. To try to ward off confusion, let me reserve the term ‘*absolutely unrestricted quantifier*’ for the (I claim) chimerical one that ranges over ‘absolutely everything’. I shall also sometimes talk of speakers attempting to *accomplish* absolutely unrestricted quantification with their utterances. By ‘unrestricted quantifier’ I just mean one that is *syntactically* unrestricted.

Arguments like the one I shall offer have a way of leading to a standoff. The premises of an argument can always be rejected, if one is willing to pay the cost. When the conclusion appears implausible, the costs of rejecting premises may appear more affordable. Which is more implausible—which cost is worse—is often difficult to tell. Hence, a standoff ensues.

My argument is no different. There are *coherent* options for rejecting its premises. But as I shall stress when we come to this issue below, at the very least, my argument will show that the costs of doing so are very high indeed. Higher than many suspect, and I believe, too high to pay. I shall show that the cost of avoiding my conclusion is either to reject some highly plausible principles about meaning, or to reject what appear to be solid reasons to conclude a certain object exists. In the end, it seems to me, the attempt to save absolutely unrestricted quantification faces an embarrassing tension. It insists on absolutely generality for quantifiers, but at the same time is willing to exclude what appear to be perfectly good objects from their domains. Generality is bought by what appears to me to be an unnatural curtailing of the notion of object, which is not much generality at all. Furthermore, I shall argue that the cost of accepting my conclusion is not nearly so high as it might appear. I shall argue that, though the restricted nature of quantification has some significant consequences, it does not generally undermine the many uses to which we put quantifiers. Hence, I shall advocate for resolving the standoff by giving up absolutely unrestricted quantification.

The discussion of this paper will proceed in four sections. The argument against absolutely unrestricted quantification will be presented in Section (I). Section (II) will discuss the impact of this conclusion, and argue that it is not nearly so implausible as it might appear. A fully accurate statement of the position I advocate will be given at the end of this section. Some

specific objections to the argument will be addressed in Section (III). Finally, Section (IV) will examine the connection between quantification and realism, and argue for the metaphysical neutrality of the position I defend.

## **I. The Argument**

Many are inclined to reject my position out of hand, on the grounds that, as it has often been said to me, ‘obviously ‘all’ means ALL’. My reply, in similar terms, is basically to say ‘yes, but what then does ‘all’ MEAN’. Oddly enough, this exchange of trivialities does point out something important. My argument is based on some ideas about meaning. In fact, I believe it is based on some ideas about meaning that are as plausible as the conclusion is implausible. I shall begin by rehearsing them. I really do mean these claims to be weak, to the point of being truisms. In presenting them, I shall try to make clear how weak they are.

### **I.1. Meaning**

*First truism: utterances only have the meanings they do because they are interpreted as having them.* Consider an utterance: I say, ‘a chair is near the door’. Consider what you just witnessed. You witnessed an event. What happened? I produced an audible *inscription* of a sentence. It could have just as well have been written, of course. What else? Of course, lots. Various things happened in our immediate and less immediate surroundings. Moreover, I made a meaningful claim. But what did you perceive? You heard the inscription being produced. But you did not observe the meaning of the utterance. You did in some way apprehend it, if you understood me, but you did not perceive it, for people have no occult meaning-sensing faculties.

How then did you understand me? How did you apprehend the meaning? Because you understand English, you can associate the correct meaning with the utterance. This is what I shall call *interpreting*.

What the hearer can perceive is not all there is to the story, of course. Among other things, genuine utterances always go along with some communicative intentions on the parts of speakers (and in many cases, we might want to give priority to the speaker's intentions over the hearer's). More importantly, the issue is not just how hearers recognize meanings. An utterance, a production of a particular physical inscription, has a meaning at all because speakers and hearers understand it as having it, and so because it is *interpreted* as having it. (In what follows, I shall suppress speaker/hearer asymmetries and simply talk of interpretation as an activity carried out by speakers.)

Now, I believe that all that is involved in this truism is the rejection of occult meaning-producing and sensing faculties. But the term 'interpretation' has been given many interpretations, so I shall pause for a moment to mention what I do not mean. I have not made a familiar Davidsonian claim. Interpretation, to me, is an entirely *programmatic* notion. It is whatever is done in attaching meanings to utterances. It need not be radical interpretation in Davidson's [1973] sense. For all I have said, it is left open that in knowing a language, we simply come to know a range of propositions, meanings, and some rules for associating them with utterances (perhaps, for instance, along the lines of Lewis [1975]).

I do want to stress that fundamentally, interpretation in my sense is something we *do*. It is what we do in recognizing an utterance to have a particular meaning.

*Second truism: meaning determines truth conditions.* There is a way this could be taken in which I don't think it is a truism at all. If we start with some basic notion of truth or truth condition, and then use it to explain meaning, we have taken a substantial position. But all I want to say here is that whatever truth and truth conditions are to be, meaning must be such as to combine with the way things are to determine truth values for utterances. Hence, in a weak sense, meaning must determine truth conditions.<sup>1</sup> Not only does this not require any particular commitment to a view of truth, it also does not require that meanings be identical with truth conditions. Properly read, I believe this is compatible with most views of meaning.

These two truisms combine. Interpretation attaches meanings to utterances, and so must fix the truth conditions of utterances.

## **I.2. The Interpretation of Quantifiers**

Now, I want to put these platitudes to work to understand utterances of quantified sentences. Say an utterance of a sentence containing a quantifier is one that *involves* a quantifier. The question I propose to ask is what is required of interpretation for it to fix truth conditions for utterances involving quantifiers. And I shall give a fairly non-controversial, but I believe important answer.

*Third truism: interpretation must provide a domain of quantification.* I do not yet claim *the domain cannot be absolutely everything*, only that some domain must be provided.

Let us focus on the universal quantifier. To interpret a sentence ' $\forall xFx$ ', interpretation must fix its truth conditions, and the question is how that is done. The simple answer is interpretation provides a domain,  $D$ , and the sentence is true if ' $Fx$ ' is satisfied by each object in

*D*. We can give other characterizations, but the observation that they make use of some domain holds for all of them. (I shall return to this point in Section (III).)

To stress, I do not yet claim that *D* cannot be ‘absolutely everything’. I have set up the background for the argument. Utterances involving quantifiers get their meanings by interpretation, and even for unrestricted quantifiers, interpretation of a quantifier requires providing the domain over which it ranges. Interpretation, as I stressed, is something done by speakers. Thus, the question of whether absolutely unrestricted quantification can be accomplished comes down to the question of what is required for a *speaker* to provide a domain, and whether it can be done for a domain of ‘absolutely everything’. The point of the argument so far has been to show that this is the question we really need to answer.

### **I.3. Domains for Unrestricted Quantifiers**

So, let us try to answer it. Is it possible for speakers to specify the domain of ‘absolutely everything’? It might, again, seem that this is a trivial matter. Cannot they do so just by saying ‘everything’? This answer will not do, for what they specify when they use this expression is precisely what is in question. But if this does not suffice, what can?

The usual way we specify domains of quantification is by using predicates. The domain is the extension of the predicate: we quantify over *people* or *stars in the heavens*. But this is of little help. Such domains usually interpret *restricted* quantifiers. Typically, these domains are given as subdomains of some background domain. People are a subdomain of living things, for example. Call these kinds of domains *derived domains* (derived, that is, from the background domain). What about a domain for a genuine attempt at absolutely unrestricted quantification?

Typical of such a domain, at least, is that it should *not* be given as a subdomain of some other domain. To introduce some more terminology, call such a domain *fundamental*.

If domains for unrestricted quantifiers need to be fundamental rather than derived, how can we specify them at all? How can we give a domain other than by using a predicate? For guidance, let us look at why a predicate does the job of specifying a domain for a restricted quantifier. The specification of a domain is part of interpretation, and so it must help to fix truth conditions. This is precisely what the predicate does. Consider a typical case of restricted quantification, e.g. ‘every book has a cover’ ‘ $\forall x(Bx \rightarrow Cx)$ ’. This is interpreted as true if and only if each member of the background domain that is a book (satisfies ‘*B*’) also has a cover (satisfies ‘*C*’). The contribution ‘book’ (‘*B*’) makes here is to divide the background domain into two parts: those objects that fall under the derived domain, and those that do not. This is the contribution we need the domain to make to the interpretation of the utterance.

It would be nice to say that we should ask the same of a specification of a fundamental domain, but of course, this makes no sense. There is no background domain to divide. But we need to ask something. Remember, the issue here is not what objects there are so much as how we can specify them in interpretation—how an interpreter can associate the right domain to a quantifier. It is clear that many specifications are inadequate for fixing truth conditions. Pointing off into the distance and saying ‘those’ will not suffice, nor will specifying a few objects, but not all of them, by saying, ‘the domain includes this, that,...’. Neither of these specify enough to determine truth conditions. The first does not fix precisely what any of the instances for an utterance involving a quantifier are to be. Hence, we cannot even use the domain to fix conditions under which an utterance is false. The second goes further. At least, a

universal quantification over the domain so specified is false if it is false of one of the objects mentioned, but the truth conditions of the utterance are still not determined, as it has not been fixed what counts as all instances, so it is not fixed what makes it true.

Each of these specifications fails to do what a predicate does with a background domain. The first, vaguely gesturing, fails to tell us what falls under the domain. The second, giving a partial list, fails in a way analogous to failing to tell us what is not in the domain. It cannot literally do this, as for a fundamental domain, there is no background domain including things not in it. But it fails to *completely* tell us what is in the domain, so we cannot use the specification to get answers to questions about the extent of the domain.

The first kind of failure, exhibited by vaguely gesturing, is a failure of *sharpness*. It is the failure to sharply tell us that certain things are in the domain. The second kind of failure, exhibited by the partial list, is a failure of *exhaustiveness*. It is the failure to tell us completely what is in the domain. Like specifications of derived domains by predicates, the specifications that suffice for interpreting quantifiers must be sharp and exhaustive. A specification that has these two properties will make the needed contribution to fixing the truth conditions of an utterance. Let me stress, it is the specifications that need to be sharp and exhaustive—objects can be neither.

How could a specification of a fundamental domain live up to these restrictions? This is a difficult question, and I will not be able to answer it fully. There is one more or less clear case, though. Suppose we assume the natural numbers form a fundamental domain, as many positions in the philosophy of mathematics do. It appears this domain is specified by a kind of generation procedure, embodied in the following inductive definition:

N1. 0 is a number.

N2. If  $n$  is a number, so is  $Sn$  (the successor of  $n$ , i.e.  $n+1$ ).

N3. The natural numbers are the closure of this process.

Such a definition cannot be given *ex nihilo*. This one relies on our grasp of the specific number 0, the successor operation, and the idea of iterating an operation to closure. There is some well-known impredicativity in such a definition, as the notions of iteration to closure and natural number are intimately related.

But though the definition itself relies on this apparatus, it does appear to specify the domain of natural numbers without simply starting with a wider domain and cutting them out of it with a predicate. (Even if they are part of a larger domain, that is not how this specification gives them.) This definition does tell us enough to fix truth conditions for quantifiers over numbers. The instances involved are precisely the things gotten from 0 by iterating the successor. It is clause (N3) that ensures the specification is exhaustive. Without it, the specification only partially fixes truth conditions for utterances involving universal quantifiers over numbers. Clauses (N1) and (N2) together ensure the specification is sharp, as they tell us that certain objects are in the domain. Together, (N1)–(N3) do just what a predicate does for a background domain. They fix sharply and exhaustively what falls under the domain, and this is precisely what we need to interpret an utterance involving a quantifier.

This goes some way towards showing that fundamental domains can be specified in interpretation. The task is not impossible. But it does not show how the fundamental domains that should interpret our widest quantifiers may be specified. My own belief is that the interpretations of our widest quantifiers is a heavily context-dependent matter, and that speakers

look to a broad range of contextual clues, both from the environment and from previous discourse, to interpret these quantifiers. But I shall not be able to investigate this matter here.<sup>2</sup>

Call a specification that is both sharp and exhaustive *determinate*. I claim, so far, that to interpret an utterance involving a quantifier, a domain for the quantifier must be specified, and furthermore, the domain must be specified determinately. Let me pause to note, providing a determinate specification in no way requires speakers to have discriminating knowledge of each of the elements of a specified domain. They only must provide some information that *fixes* determinately what is in it. For our widest, unrestricted, quantifiers, the domain specified must be a fundamental domain. I have noted one example of how such a domain might be specified determinately.

Now, we finally get to the real issue. I now must show that there cannot be a determinate specification of a domain of ‘absolutely everything’. This will show that our quantifiers cannot range over such a domain.

It is tempting to suppose that in fact, speakers can quite easily give a determinate specification of such a domain. They may simply employ a predicate like ‘object’, and indeed it seems they need only say ‘the maximal domain is to consist of *all objects*’. Those unhappy with the term ‘object’ may prefer circumlocutions like ‘the maximal domain consists of all  $x$  such that  $x=x$ ’.

Observe, however, that for this suggestion to get off the ground, it must place a great deal of weight on the predicate ‘object’. The second formulation, for instance, threatens to collapse into circularity, if it relies on the unrestricted universal quantifier to specify the domain of the unrestricted universal quantifier. To get a non-circular specification, the work of specifying the

domain must be done by the concept of *object* expressed by the predicate, rather than the universal quantifier.

Does this concept of object suffice to give determinate specification of a domain of absolutely everything? I shall argue it does not, in two steps. First, I shall observe that our ordinary talk using the term 'object' is much too vague to yield a determinate specification. This much should come as no surprise. But second, I shall argue that this feature of the ordinary term cannot be overcome in a way that could suffice to capture 'absolutely everything'. In particular, I shall argue, ways of making the notion of object satisfy the sharpness requirement must fail to be exhaustive, or fail to be inclusive enough to be reasonable candidates for 'absolutely everything'. This will show there can be no determinate specification of absolutely everything.

It may go without saying that our ordinary concept of object, which we invoke when we simply make use of the term 'object' of English, is hopelessly vague. One feature of this will be important in what follows. Consider one of our students in Introduction to Philosophy, who simply rejects out of hand the existence of highly scattered objects (mereological sums of objects highly scattered in space and time). Suppose this student simply insists that such things are not objects. This student begs the question about the ontology of material objects, no doubt. But not in a way that is revealed by our ordinary concept of object. *If* the student is mistaken, it is not a matter of her being somehow incompetent in the English language, misunderstanding the term 'object'; nor will any appeal to this concept help show the student her mistake. We need not look to beginning students for these sorts of examples. Current defenders of Meinongian universes containing 'non-existent objects' may be accused of many errors, conceptual or factual, but it would be unfair to accuse them of incompetence with the English language. The same may

be said of historical figures who held it to be (practically) analytic that all objects are concrete (though Latin or German may be the more relevant language).

The important feature here is not simply that the term ‘object’ is vague. It is, but that simply means it may need to be sharpened to provide a sharp specification. (I introduced ‘sharpness’ as a term of art, but it is clear enough that genuine vagueness precludes a predicate’s providing a sharp specification.) What then becomes important is that the term itself is unable to decide amongst its potential sharpenings. The student is intuitively relying on a particular (partial) sharpening of the notion of object, and so finds it trivial to say there are no scattered objects. There may be good reason, both factual and conceptual, to reject this sharpening as inadequate, but these are not provided by the meaning of ‘object’ itself. Likewise, a sophisticated Meinongian will in part rely upon some (partial) sharpening of the notion of object, as well as arguments about what falls under it. Again, the ordinary meaning of the term ‘object’ is not sufficient to decide whether the sharpening involved is preferable to another, more limited one. We will not refute such a philosopher by rough-and-ready lexicography.

Not deciding amongst its potential sharpenings is a general feature of vague terms. Consider, for instance, the predicate ‘heavy’. We can imagine many reasons to sharpen this predicate. Perhaps we are trying to write industrial rules about how heavy a load someone is allowed to carry, for example. But for any reasonable sharpening we choose, there will be another more inclusive sharpening. If 70lb. is too heavy to carry, we could just as well say 69.9lb. is. One could debate the merits of various regulations sharpening ‘heavy’. One might hold that 69.9lb. is not too heavy to carry, for whatever reason. But (just as with our introductory student), one cannot say that 69.9lb. is not heavy *because* 70lb. is the boundary for heaviness. A

vague predicate does not itself decide among its potential sharpenings. It is doubtful that this is all there is to vagueness, but it is an important aspect of it.<sup>3</sup>

This vagueness-like feature of the ordinary term ‘object’ shows why appeal to this term, by itself, is of little help in providing a determinate extension of ‘absolutely everything’. Not only does the ordinary term fail to provide a determinate specification, it also fails in any direct way to guide us towards one. We need to sharpen the meaning of the term, to get a determinate specification of anything. But if we attempt this, we run the risk of hitting upon too narrow a notion of object to really give us absolutely everything. We are unlikely to make the particular mistakes I mentioned above. Few, even nominalists, would accept that it is a feature of the concept of object that objects are concrete. Likewise, few even among those who worry about scattered objects, would be happy with a notion of object that rules them out by itself. But as we have to sharpen our ordinary concept of object somehow, we do run the risk of sharpening in such a way as to entail some more recherche restriction. We run the risk of hitting upon a sharpened conception of object for which there is a nearby sharpening which is more inclusive. And as I have stressed, we cannot look to the ordinary meaning of the term ‘object’ to guide us. It does not decide amongst its potential sharpenings, so it will not help to avoid overly narrow ones. If we are to have a sharpening which could plausibly give a determinate specification of ‘absolutely everything’, we need to be sure it is not so overly narrow. Thus, we cannot rely upon our ordinary term ‘object’ to do this task.

I have so far argued that we must provide a determinate specification of a domain of quantification, and that appeal to the ordinary term ‘object’ will not help to provide a determinate

specification of a domain of absolutely everything. This does not yet show there cannot be such a specification. To see if there can, we need to consider other ways it might be provided.

In specifying such a domain, we are not seeking to answer all ontological questions, only to find some way of fixing meaning that will work. We thus want the minimally detailed, but maximally broad specification of a domain of objects. And indeed, a certain amount of progress has been made in articulating such a specification. It has proved fruitful to pursue this project with the aid of logic. By asking what the logical characteristics of objecthood are, we might hope to delineate with the precision that logic brings the outer limits of our notion of object. We might hope to give determinately the *maximally broad conception of object*, and use it to interpret unrestricted quantifiers. This should alleviate the worry I pressed about accidentally stumbling into unduly restricted sharpenings, as the hope is to use logic to build a sharp specification from the ground up that builds in maximality.

Looking for a maximally broad conception of object in logic has an admirable pedigree, going back at least to Frege.<sup>4</sup> I am going to skip the character buildup, though, and cut right to the chase. As quantification is at issue, we can expect the conception to be articulated by some characterization of the logical notion of singular term. The general, *logical notion of object* may be taken to be that of to what a singular term, so understood, refers. There are some well-known and important worries about how and whether this can be done, and some of them are especially pressing given that quantification is at issue.<sup>5</sup> However, I think supposing there is such a characterization gives proponents of absolutely unrestricted quantification their best case, so I only make my argument stronger by not worrying about them here.

The question now becomes how some such notion can give a *determinate* specification of a domain of quantification. It might be able to, as follows. We are granting a determinate specification of what it is to be a singular term. We then say that the objects are all and only the potential referents of singular terms. Objects, roughly are the kinds of things that are named by terms. Again, there are many points here that may be challenged, but I think I offer the defender of absolutely unrestricted quantification the best case by letting them go. So I think it is fair enough to grant that such a specification is at least sharp.

But there are difficulties we cannot ignore. For such a specification to be coherent, the notions of singular term and reference must be prior to the interpretation of unrestricted quantifiers. These notions are used in specifying the interpretation. For the same reason, the notion of object in general, the logical notion, is prior to and independent of any specification of a domain of quantification. These priorities have a consequence that will be of great importance in a moment. Because of them, we have no prior guarantee that the referent of any particular term will be in a particular specified domain.

The result is that though we may grant that we can use the logical notion of object to produce sharp specifications of domains, exhaustiveness and maximality are another matter. Though the specification is certainly general enough to avoid the kinds of restrictions I worried about a moment ago, we have not built in any guarantee of maximality that insures we will get a sharp and exhaustive specification of absolutely everything. We can still discover, once we get a sharp and exhaustive specification, that we can also find a term whose referent falls out of the domain so specified.

In fact, this is what happens. Under the logical notion of object, it is impossible to give an exhaustive specification that encompasses the entirety of the logical notion. More precisely, for any determinate specification, we will be able to name at least one object not falling under it, by appeal to the logical notion of object.

This may be shown in several ways. One is by way of Russell's paradox. Suppose we have some specification of a domain. As a result, we can quantify over it. Hence, we can form the class term ' $\{x:x=x\}$ '. (A class term can be thought of as being given by the stipulation that  $\{x:\phi(x)\}=\iota y(\forall z(z\in y \leftrightarrow \phi(z)))$ .) As a well-known corollary to Russell's paradox, this cannot be in the domain over which ' $x$ ' ranges. Suppose  $y=\{x:x=x\}$  is in the domain. Then by comprehension (in fact, restricted comprehension), there is a set  $z=\{x\in y:\neg x\in x\}$ . But  $z\in z \leftrightarrow \neg z\in z$ , so we have a contradiction. We thus have a term that cannot refer to anything in the specified domain.

This version of Russell's paradox is simple, and familiar. But it needs to be stressed, the appeal to class abstracts, or sets, or such things, is not crucial to the argument. Rather, the point is that the process of interpretation, particularly the interpretation quantifiers, suffices to reveal an object that falls outside of the domain of that interpretation. The domain of quantification itself is such an object. To make this more vivid, let us consider a somewhat more general version of the Russell argument, due to Timothy Williamson [MS]. Williamson invites us to consider an interpretation  $I$ . In the terms of the discussion here, we may think of  $I$  as the product of the process of interpretation undertaken by speakers. As Williamson notes, we need to say little about the formal properties of  $I$ , except that for a given predicate term ' $P$ ' and collection  $F$ , we may build an interpretation  $I(F)$  which makes ' $P$ ' hold of all and only the  $F$ s.  $F$  need not be a

set or class. Those fond of the plural reading of second-order logic (as in Boolos [1984]), for instance, may simply take the process of interpretation to specify that ' $P$ ' holds of all and only the  $F$ s without requiring  $F$  to be a class.

This is already enough to generate a version of Russell's paradox. Interpretations need not be sets, but they are objects. Hence, we may let the  $R$ s be all and only the objects  $o$  such that  $o$  is not an interpretation under which ' $P$ ' applies to  $o$ . Then there is an interpretation  $I(R)$ . But  $I(R)$  cannot be in any quantifier domain given by  $I(R)$  itself. If it were, we would have that according to  $I(R)$ , in the case of  $o=I(R)$ :

' $P$ ' applies to  $o$  iff ' $P$ ' does not apply to  $o$ .

This is a contradiction, just as in the class-based Russell argument.<sup>6</sup>

This version of Russell's paradox shows that there is nothing specific about classes or sets involved; rather it is the role of interpretation that is crucial. To make this role clear, let me pause to review the structure of the case that has been offered against absolutely unrestricted quantification. I first observed that interpretation is essentially involved, and in particular, observed that interpretation of even a syntactically unrestricted quantifier requires providing a determinate specification of a domain of quantification. I then asked whether such a specification of 'absolutely everything' could be given. I argued that a plausible attempt must be based the logical notion of object, which seeks to capture the maximally broad conception of object. I went on to argue that even this must fail to give a specification which is at the same time sharp, exhaustive, and plausibly absolutely everything. It is here that the paradox enters the picture. It shows that if we have a determinate specification of a plausible candidate domain, we can use it to find another object, which cannot on pain of contradiction be in the domain

specified. Thus, the paradox shows that plausible determinately specified domains cannot be absolutely everything.

The objects the paradox leads us to are what we might think of as ‘artifacts of the process of interpretation’. The domain of a sufficiently wide-ranging quantifier itself is an object that cannot fall within the domain of the quantifier. This is for Russellian reasons, but it is not basically a point about sets. Rather, it is the process of interpretation that is important. This is highlighted by the more general version of the paradox, which talks only about interpretations themselves. The process of interpretation, especially the interpretation of wide-ranging quantifiers, itself generates objects—or perhaps better, instructions for *identifying* objects—which must fall outside any domain given by the interpretation. The paradox shows that the process of interpretation can, in extreme settings, lead us out of any domain of the interpretation we produce.

#### **I.4. The Lesson of the Paradox**

As I noted at the very outset of this paper, the results of an argument like the one I have just offered are never entirely straightforward. As this one results in a conclusion many will find implausible, it invites rejection of one of its premises, or one of its steps. We should consider what might be done to avoid the conclusion of the argument I have offered here, and what the costs might be. This discussion will also help clarify the role of the logical notion of object in the argument.

Which premises or steps might we plausibly reject? I tried to put the initial premises about interpretation in such general ways as to make them appear as truisms. At least, I do not

see much prospect for arguing against them. Rather, I think the most likely place to resist the conclusion is at the point where Russell's paradox enters. There, one might try to argue that there is no such object as the one I used the paradox to identify. As this object demonstrated that a plausible candidate for a domain of absolutely everything could not be one, rejecting the existence of the object allows one to maintain that the domain really is absolutely everything.

In the class-based version of Russell's paradox, for instance, it might be tempting to say there is simply no such thing as the universal class. However, I think that the appeal to interpretation, and the more general version of Russell's paradox, make it difficult to reject this part of the argument. We have already observed that the conditions on interpretations are highly general. Hence, no specific ontological worry about sets or classes will suffice to reject the existence of the object in question. More importantly, I think we simply do have good reason to grant that there is such an object. After all, the process of interpretation itself shows us what the object is. Let us talk about this in class terms. We are looking at an object like  $\{x:x=x\}$ , where 'x' ranges over the elements of the domain specified by an interpretation. Insofar as we really have given a determinate specification, we have specified exactly what falls in this class. As classes are extensional, I see nothing else that could be required to convince ourselves that this object exists. Without appealing to classes explicitly, we may observe something similar. The process of successful interpretation shows us an object which cannot be in the domain of the interpretation. Insofar as we have been successful in interpretation, we have carried out the process of identifying the object. Again, I do not see what more could be required to show something exists. I must grant that denying its existence is coherent. But outside of the

question-begging insistence that the original domain was absolutely everything, and hence there can be no such object, it is hard to see why a plausible case for its existence has not been made.

The question at hand can be well-framed by the logical notion of object, to which I said we must appeal to provide a maximally broad determinate specification of a domain. This conception of object says that objects are the things to which singular terms refer. Hence, to determine whether we have an object, we have to decide whether we have a well-formed singular term, and whether or not it refers. So far, I have taken for granted that we have a well-formed singular term, and I have argued that we have good reason to then conclude it refers. If so, then according to the logical notion of object, nothing more need be said.

Another way for the argument I have offered to be resisted, while still accepting the logical notion of object, is to challenge our having a genuine singular term. Though I do not think this option is acceptable, it seems (somewhat surprisingly) to better capture the basis of my opponent's objection. I have argued that the interpretation speakers construct itself provides all the evidence we need that the disputed object exists, by showing how to identify it. In insisting there is still no such object, I think my objector is insisting that there is something else that is needed. Under the logical notion of object, the only way there could be is if the status of the singular terms provided by the interpretation is questionable.

To give this more vividness, it might be helpful to return to the idea of the plural interpretation of second order logic. This might invite one to agree that I have indeed specified some objects, the ones falling under the domain (or in the more general version, those that are somehow involved in the process of interpretation). But such an objector might still insist I have not identified an object, on the grounds that I have not provided a well-defined singular term

which picks out that totality. At best, the objector might grant, I have given a plural term, which is not singular, and so not a guide to objecthood.

I do not think the plurals response is correct here, but it helps to make vivid the basic dispute. I do not think the plurals response is right, as insofar as our grasp of the plural idioms used to interpret second-order variables derives from natural language, it equally shows us how to find what are clearly singular terms.<sup>7</sup> Second-order variables range over predicate positions. But in English, we can easily *nominalize* these, producing singular terms. A predicate like ‘red’, for instance, is nominalized as ‘redness’. The plural reading invites us to use the apparatus of natural language plurals, but the very same natural language provides us with nominalizations which give us perfectly good singular terms. Natural language, being rather insensitive to delicate ontological distinctions, may equivocate between ‘the group/collection/set of things falling in the domain’. But regardless, it does provide us perfectly good singular terms. Thus, I think the attempt to resist my conclusion this way is not viable either.

The basic tension here, I believe, is between the evident ability of natural language to make extremely wide generalizations, and its evident ability to make what we might call ‘widespread nominalization’. Linguistically, this is the ability to introduce singular terms for many elements that do not originally appear in nominal positions. In English, as I observed, we can nominalize predicates. But using pronouns or definite descriptions, our ability to nominalize in English appears much greater. To mention one more example, consider:

John walked quickly. The quickness with which he walked was great. It was faster than usual.

We do not really need the full nominalizing power of English to argue against unrestricted quantification. What we need is to be able to nominalize the process of interpretation, as the more general version of the Russell argument of Section (I.3) showed. This gives us objects which cannot fall within the domain of the unrestricted quantifiers of the interpretation, even if the domain endeavored to be maximal.

Of course, it is still coherent to insist that this widespread pattern of nominalization in natural language is merely a misleading surface feature, which leads us to mistakes about genuine logical form. This would allow one to insist we do not have genuine singular terms, and so we do not have an object at the crucial point in my argument. However, I believe the cost of doing this is very high indeed. The appeal here is not to some special theoretical entities, such as sets or class; it is merely to the guidance we get from simple linguistic usage, clear to any competent speaker. Any speaker sufficiently reflective about the process of interpretation can step back and say 'it'. To reject the nominalization we need to find objects outside of the quantifier domain, we need to declare this incoherent.<sup>8</sup>

Insofar as my opponent does have a coherent position here, as I promised at the outset, we have reached a standoff. Both sides in the standoff have some embarrassments. I have to insist that apparently unrestricted quantification is not really so. My opponent, on the other hand, has to insist that apparently successful nominalization is not so. I have tried to show that my opponent's position is unattractive. Insofar as we are accepting the logical notion of object, my opponent must either claim the term in question is not a genuine singular term, or it does not refer. I argued that a successful process of interpretation provides good evidence the term does refer. It thus seemed to me the most promising route was to insist the term involved is not really

singular, and hence, to put it intuitively, ‘we don’t really have an object’. But I pointed out that the common natural-language phenomenon of nominalization gives us good reason to think we do have a perfectly good singular term. My opponent, just as much as me, has to deny what appears to be an obvious feature of ordinary usage.<sup>9</sup>

If I am right, then we are precluded from achieving absolutely unrestricted quantification. At best, absolutely unrestricted quantification can be defended only by sacrificing the equally compelling widespread nominalization. Once we see the role of interpretation in understanding quantifiers, the pressure to nominalize, I have suggested, is all but irresistible.

The result is that the logical notion of object, by which we sought to articulate a maximally broad domain, displays a strong form of the feature related to vagueness we saw above. In particular, any determinate specification of a domain of quantification itself provides the resources for identifying an object not in that domain, and hence for specifying a strictly broader domain. The logical notion of object thus exhibits what Dummett [1991a, 1993] has called *indefinite extensibility*. Insofar as the logical notion of object can be used to produce a determinate specification of a domain of quantification, it produces not just one, but an indefinitely increasing sequence of them. We may think of each determinate specification as an articulation of the logical notion of object. But the generality of the logical notion is inconsistent with any such articulation both being determinate and capturing all of the logical notion of object. We may move from any such articulation to a broader one, still an articulation of the logical notion of object. Hence, no determinate specification gives ‘absolutely everything’.<sup>10</sup>

## II. What Can We Say?

This concludes the argument that absolutely unrestricted quantification is impossible. If this argument is right, what are its consequences? Is its conclusion really so counter-intuitive as it might seem? I do not think it is. In this section, I shall discuss why.

First of all, let me repeat that the conclusion does not show that there are certain things we cannot quantify over. I have not said, for example, that we cannot quantify over all sets.<sup>11</sup> We *only* can never quite get absolutely everything.

Furthermore, I believe that the linguistic intuition that we simply do accomplish absolutely unrestricted quantification is less compelling than it might seem. In fact, my position has little impact on most discourse. Two striking facts about natural language support this. First, most, if not all, quantification in natural language is *syntactically* restricted. We say ‘every book’, or ‘everyone’. We can, of course, say ‘everything’, but this appears to be syntactically complex, containing a restrictor.<sup>12</sup>

Second, and perhaps more importantly, semantically even the expression ‘everything’ is rarely used with great generality. Someone who says ‘Everything was destroyed by the fire’ does not speak of a great ontological cataclysm, for instance. Natural language quantifiers are context sensitive. Particularly, the domain over which a quantifier ranges in a given usage is determined in part by context. This is especially so for constructions like ‘everything’, in which the nominal ‘thing’ does little to fix the domain of the quantifier. Aside from the ‘fire’ case, I might mention a cookbook on my shelf entitled *How to Cook Everything*.<sup>13</sup>

In most of these cases, we see quantifiers like ‘everything’ used as semantically restricted quantifiers, in that their contextually determined domains are subdomains the background

domain provided by the context. What the extensibility of the logical notion of object shows, I suggest, is that even the background domain of a context—the local maximal domain over which the widest quantifiers used in the context range—is apt to expand. Thus, there will be other contexts in which a wider domain is available. The paradoxical reasoning we rehearsed in Section (I.3) shows one way to induce this kind of context shift. Thus, in a sense, all quantifiers are contextually restricted, even the widest quantifiers of a context. But it must be stressed, this is not to say that all quantifiers are restricted in the sense of ranging over some subdomain of the background domain of the context. We can use quantifiers to range over the background domain, of course. But there are other contexts in which even this domain can be expanded.

The processes by which the background domain of a context is set need to be investigated, and I shall not have the space to do so here.<sup>14</sup> But at least, we can observe that reasonable interpreters in reasonably normal contexts will set their background domains to be quite large. They will include, let us suppose, all the medium-sized dry goods and most abstract or theoretical entities. Given this, the arguments of the last section show only a very special sort of expansion. The objects discovered in the paradoxical reasoning are generated by the interpretation of the unrestricted quantifiers themselves. As I put it above, they are artifacts of the discourse. The expansion of background domains observed is simply by adding these artifacts, generating a new wider context. But this expansion will not affect the domains of medium-size dry goods, or objects we learn of from fundamental physics, or numbers, etc. It is hard to see much practical impact of this sort of context shift on most of what we say.

The result, according to my view, is that there are limits on how interpretation works, but they have little impact on most of our speech. Yet there are, no doubt, a few cases where we

really strive to accomplish absolutely unrestricted quantification, and will not be satisfied with a quantifier ranging over even an artfully fixed but contextually restricted background domain. Examples include logical truths, such as ‘all objects are self-identical’ ( $\forall x(x=x)$ ). For these, and some other cases, extraordinary measures may be required.

Let us look at the logical truth case a little more closely. Part of what makes a logical truth a logical truth is that it is insensitive to features of domains of quantification, and I believe this is what makes the example so compelling. It is so plausibly about absolutely everything precisely because it does not matter over what its quantifier ranges. But once we see this, we can see that, though it must be interpreted as having its quantifier range over some particular domain, it can also be construed as offering more information. It can also tell us, ambiguously, a feature of any domain it might be interpreted as ranging over.

This is to propose to treat logical validities as typically ambiguous, in something like Russell’s sense. This idea has already been much-discussed, so I shall limit myself to a few observations.<sup>15</sup> Making and understanding a typically ambiguous utterance requires more from speaker and hearer—more from interpretation—than normal utterance does. It requires, over and above interpretation, the observation that the truth of the utterance does not depend on what the domain is. This seems possible for some claims, like logical validities, but there are limits to how widely it can be done. The range of potentially ambiguous utterances is probably wider than logical truths, though. What needs to be established is not so much independence of truth from the domain entirely, but something like what logicians call persistence: insensitivity to expansion of the domain.

For an utterance to be typically ambiguous is for its interpretation to bring with it some extra information. Like recognizing meaning, recognizing typical ambiguity is a skill that can be acquired by speakers. But we should *not* see typical ambiguity as a special kind of meaning, and so we should not see typically ambiguous utterances as accomplishing absolutely unrestricted quantification in anything but a roundabout way. The meaning of a typically ambiguous utterance is still fixed by its interpretation, with whatever domain that specifies. But with a typically ambiguous utterance, once its meaning is identified, it is possible to go beyond interpretation, by recognizing that the utterance's truth conditions are insensitive to expansion of the background domain of the interpretation. The utterance can thus be thought of as holding ambiguously as the domain might expand. Hence, typical ambiguity is an extraordinary measure, over and above ordinary interpretation, which may work in some special cases but not others.<sup>16</sup>

Now, I can finally explain the proviso to my position, and give it a full and proper formulation. For any domain of quantification that can be specified in interpretation, there can be specified a wider one. Hence, no utterance can accomplish absolutely unrestricted quantification, in the sense of being interpreted as having its quantifiers ranging over a maximal and non-extensible domain of 'absolutely everything'. In some special cases, including logical validities, interpretation can be supplemented in such a way as to allow speakers to recognize utterances as typically ambiguous, and so carrying additional force. This explains the more compelling of the chimerical appearances of 'absolutely unrestricted' quantification.

Does the combination of contextually determined background domains and typical ambiguity undercut the force of my conclusion? It is certainly mitigated, but I believe it is still quite substantial. Interpretation is always relative to a contextually determined background

domain: the local maximal domain over which unrestricted quantifiers are interpreted as ranging when context does not further restrict. This domain, if I am right, is apt for expansion. If we encounter paradoxical discourse, like that evident in the Russell's paradox reasoning I recited, or in discourse related to the Liar paradox, the domain may expand as context shifts during the discourse. I am inclined to see this as an insight rather than a problem. It helps to highlight a genuine feature of discourse. To borrow a phrase from Kripke [1975], speaking and interpreting is a *risky* enterprise. Risky in many ways. Some are mundane. We risk misunderstanding, or being misunderstood, or even just being inadvertently inappropriate. But risk has other aspects as well. In some extreme situation, we risk failing to capture quite all of what we might have been trying to talk about. In paradoxical cases, where logic steps in, our attempts to make extremely broad generalizations can go wrong. We can find that we left something out, as the argument above shows.

We have seen this risk can be mitigated in several ways. One can to some extent insure against the risk, by making the initial contextually determined background domain large enough, and relying upon the fact that the paradoxical arguments only add what I labeled 'artifacts of the discourse' to the domain. This will ensure that most restricted quantifiers are untouched by the expansion. One might also attempt to overcome the risk, in some special cases, by the observation of typical ambiguity. But the risk is still present.

### **III. Objections and Replies**

Now that we have a full and proper formulation of my position, I shall consider some common objections. The first relates to Section (I), but is sufficiently technical as to deserve treatment outside of the main line of argument.

*Objection:* In the case of unrestricted quantifiers, the rules governing the quantifiers themselves suffice for interpretation, and so no specification of a domain is needed.

*Reply:* This simply not so. This was pointed out above for a basically Tarskian treatment of quantifiers. It is equally obvious for generalized quantifier treatments, now common in linguistics. Take, for instance, a sentence ‘ $\forall xFx$ ’. The generalized quantifier approach takes it for granted that the universal quantifier ranges over a set **A**, and that the extension of ‘*F*’ is also a set **B**. Then ‘ $\forall xFx$ ’ is true just in case  $|\mathbf{A} \setminus \mathbf{B}| = 0$ .

A more prosing approach for the objection is a syntactic characterization. Consider again ‘ $\forall xFx$ ’, supposing ‘*x*’ to be the only free variable in ‘*F*’. Syntactically, the universal quantifier is characterized by the two rules:

$$\forall\text{-Introduction: } \frac{\Delta}{F} \qquad \forall\text{-Elimination: } \frac{\Delta}{\frac{\forall xFx}{F[x/t]}}$$

with the usual provisos that for the introduction rule  $\Delta$  does not contain any undischarged assumptions containing ‘*x*’ free, and in the elimination rule ‘*t*’ must be substitutable for ‘*x*’.

In this case, though the rules themselves do not require that a domain be specified, their application in interpreting utterances does. To see this, let us look at the rules a little more carefully. They can be thought of as rules for assertion or for determining truth. We can assert a universally quantified sentence (it will be true) if we can assert the formula (it will be true) that

results from dropping the quantifier and leaving the previously bound variables free, provided our reasons for asserting the latter (the grounds for its truth) do not trade on the particular value that the free variable might be assigned. (This is the proviso that  $\Delta$  cannot have any undischarged assumptions containing ' $x$ ' free.) The point of the rule is that we can assert the quantified sentence if we can assert that  $F$  holds of an *arbitrary* object  $x$ , which is assured by finding grounds for asserting ' $Fx$ ' not relying on any assumption that might tell us what  $x$  is supposed to be. It is the notions of bound and free variables that captures the needed aspects of arbitrariness. Free variables are referential terms, but their reference is allowed to be determined arbitrarily in interpreting a formula. So if we are justified in asserting ' $Fx$ ' with ' $x$ ' free and nothing in  $\Delta$  constraining its interpretation, we can conclude that  $F$  holds of an arbitrary object.

Because they rely on some notion of arbitrary object, these rules can only fix truth values with the help of a domain of quantification. There is, of course, no separate category of arbitrary objects, mysteriously different from specific objects. Rather, the issue is *arbitrary reference*. Arbitrary reference only makes sense if we are given a domain of objects, and the stipulation that a term refers to some unspecified object in that domain. This is brought out most clearly in how we might apply the rules to determine a truth value. We can conclude that ' $\forall xFx$ ' is true if we can conclude that ' $Fx$ ' is true for arbitrary  $x$ . Yet we could only conclude this on the basis of some specification of the domain of things  $x$  can be. The same holds for the elimination rule. To apply it, we need to know from what class we may draw singular terms. If the language is interpreted, this requires knowing over what domain the variable to be instantiated ranges.

*Objection:* The position is self-defeating, as it is necessary to talk about absolutely everything to state it.

*Reply:* First of all, it is not obvious this is so. Interpretation, as I stressed above, is an activity—it is something we do. Insofar as my position is one about the nature of interpretation, it is hardly one that attempts to talk about ‘absolutely everything’, but only a specific range of human activities. I have argued that, especially once we appeal to the logical notion of object, we find that certain aspects of the interpretive process correspond to objects. Interpretations themselves, or the domains of quantification they specify, are objects. But here to, we are not strictly speaking talking about absolutely everything, but rather about a highly restricted class of objects that correspond to aspects of the process of interpretation.

To be fair, there is a sense in which more general talk is *implicit* in talk about domains. If we could ‘read through’ talk about domains to talk about the objects in them, then we might see talk about all domains or all interpretations as talk about whatever falls under them. But note, to do so is not simply to attach the usual meanings to the utterances in question, but to see something extra about them. This is just what I see happening in cases of typically ambiguity, and this case is ripe for such treatment. The discussion here depends only minimally on the specifics of what is in any domain. I sometimes have supposed that certain domains include certain objects: specific classes, or interpretations themselves. Other than that, no claim I have made depends in any way in what falls under any domain under discussion. Hence, everything I have said appears to be persistent: it is insensitive to expansion of domains. This is sufficient, I suggested, for typically ambiguous reading. More needs to be said to be certain that typical ambiguity is licensed in this case. Likewise, more needs to be said to understand just how much we can ‘read through’ talk about domains. But it does appear safe to conclude that the extent to which we can ‘read through’ talk about domains to talk about what falls under them is just the

extent to which we can see the claims I make here as typically ambiguous. Hence, whatever force there is to reading through my talk about domains is not enough to make the position I have offered self-defeating.

*Objection:* Philosophical discourse, especially ontology, requires absolutely unrestricted quantification.<sup>17</sup>

*Reply:* My reply to this objection has two parts. The first, weaker, part, considers how this sort of objection might fit into the dialectical situation. The second part looks more carefully at what my claims imply about ontological discourse.

First, to the dialectical situation. I have granted throughout that the conclusion that absolutely unrestricted quantification is impossible will strike many as counter-intuitive, if not implausible. I argued in Section (II) that it is not as implausible as it might seem, but I did grant there that it may have some unwelcome consequences nonetheless. In light of this, to point out that the state of some ontological discourse might be among the unwelcome consequences seems to me to have little argumentative force.

As I have mentioned at a few moments, we might find the arguments of Sections (I) and (II) to lead to a standoff. If the conclusion is really implausible enough, we can reject the premises or steps of the argument. But the appeal to ontological discourse does not seem to help us understand what might be wrong with the premises of the argument I offered. It is especially hard to see how it favors any of the options for rejecting the argument I considered in Section (I.4).

As I described in Section (II), one of the upshots of the argument I have offered is that extremely wide generalization can be *risky*. Rather than show us the implausibility of my

conclusion, it may be that ontological discourse highlights this risk. It may be that some ontological discourse strives for more than can be achieved. We certainly need not look far into the ontology literature to find examples of risky moves gone awry. Not only failed generalization, but all sorts of unintended content, or outright incoherence, can be the result.

The second part of my reply is to consider what possibilities may remain for ontology, given my conclusions. It is not really fair to hold the endeavor responsible for its less careful or cautious moments, so to note that it is risky and leave it is not enough.

It seems to me there are ways for ontology to go forward, even in the face of my conclusions, and even though they are risky. I shall mention two. First, it is not so clear how much of ontology really requires absolutely unrestricted quantification. Much ontological work is done by drawing distinctions among kinds of objects, or by deciding into which kinds certain ranges of object fall. Debates about the ontological status of, for instance, events or universals or mathematical objects, appear to have this character. If such talk does not really require talking about absolutely everything, but rather everything of some specific kinds, it is not threatened by the conclusions I have drawn here. Even if some ontological discourse is risky, it seems that some important parts of it are not so risky.

Second, even if some more risky ontological discourse does require greater generality, we should consider that when it is successful, it may be ripe for treatment as typically ambiguous. Particular claims would have to be examined with care, to be sure. But generally, it is sometimes the case that an interesting claim in ontology does not care what the exact extent of the domain in question is. Rather, such a claim may mark out features of the domain that are insensitive to its extent. I already mentioned that the sort of talk I entered into in describing the logical notion of

object may be treated as typically ambiguous. So would be the related Quinean dictum that ‘to be is to be the value of a variable’. It is not easy to tell just what will fall properly into this category, especially given the difficulties in establishing typical ambiguity I stressed in Section (II). But to the extent that we really do need extraordinarily wide generality in ontology, typical ambiguity gives us a way to achieve it in some special cases.

I am not sure just what claim in ontology will fall into either of these two categories. But they do suffice to show that my conclusions leave a great deal of room for ontology to go forward. The whole endeavor is by no means undone. It may be made more difficult. I do grant that the arguments I have presented may highlight a way that ontological discourse is unusually difficult, and especially risky. But, I suspect, when it comes to ontology, we knew that it was already, and the argument I have offered cannot be rejected because of that.

#### **IV. Realism**

So far, I have offered an argument against absolutely unrestricted quantification in Section (I). In Section (II), I defended my position by attempting to defuse some of the intuitions against it. In Section (III), I considered some specific objections, the last being a philosophical concern about the status of ontology.

In this section, I shall turn to another philosophical concern, which I believe needs a more extended treatment. The position I have advocated here might be taken to force upon us some sort of *anti-realism*.<sup>18</sup> Realism and anti-realism are difficult theses to state, but let me start by giving the intuitive idea. Realists hold that the world is made of objects that exist independently of ourselves, our thoughts, or our language. What there is is thus, in some metaphysical sense,

determinate, independently of us. Equally importantly, the realist must insist that it is this world of objects that our thoughts and claims are about. But if so, we might suppose, we can certainly talk about all that stuff, whatever it is, and so we can talk about absolutely everything. The only way we could not be able to, it might be suggested, is that what there is is somehow dependent on what we say or think, so that there is not a single determinate maximal domain prior to thought or language.

I think this is a mistake. The position I have defended is to a great extent metaphysically neutral. Its force is more linguistic than metaphysical. The reason absolutely unrestricted quantification cannot be accomplished is that we cannot specify absolutely unrestricted domains. This is a limitation on our ability to interpret, and so to speak and understand. As I put it a moment ago, it is a way in which interpretation is risky. But it is not a limit on what the stuff we talk about is like—whether it is in some metaphysical sense determinate. The difficulty of interpreting, difficulty driven by the combination of requirements on meaning and by paradoxical cases, may make it difficult to make claims with a certain sort of generality. This is not any kind of anti-realism.

Let me pursue this a little further. One way realism is sometimes put, in this instance by Hilary Putnam, is as follows, “On this perspective [that of the realist, or ‘metaphysical realist’ as Putnam calls it], the world consists of some fixed totality of mind-independent objects. There is exactly one true and complete description of ‘the way the world is’ ” [Putnam, 1981, p. 49]. It may be natural to read this as a claim about domains of quantification, but we must not, and it is not necessary to do so. If we understand the key terms ‘fixed’ and ‘mind-independent’, we have most of its force, and these do not require claims about quantification.

The more serious worry is that the requirement that there be ‘exactly one true description’ might generate a conflict with my view, but it does not. The threat to realism is *mutually incompatible* descriptions. For example, as Putnam moved away from realism, he began to consider how we might compare the kinds of descriptions of the world generated by different theories, or languages. In one example [Putnam 1978b, p. 131], he asked us to consider what might decide between a description of the world in which there are points on lines, and one in which there are none, but treats them as logical abstractions about convergent sets of line-segments. Putnam’s position of the time, and many more drastic anti-realist positions, hold that there is no choosing between these two descriptions. The two are incompatible, so the conclusion is to be that what the world is like is at best a theory-internal matter. The realist, on the other hand, would insist there is a fact of the matter about which one is right, depending on whether there are in fact points.

The situation we find ourselves in if my claim is correct is very different from the one anti-realists envisage. What is characteristic of the cases I pointed out, that call for extensions of domains, is precisely that what they call for are *extensions*. We can, in some difficult cases, convince ourselves we left something out of a domain. This does not call for wholesale revision of any theory. The extended domain is entirely commensurable with the previous one. It differs only in having at least one more object, and calls for no new descriptions of the objects that were in it before. Finding a new object does not require us to think that it was not there before we found it, nor that it is somehow dependent on us or our theories, nor does it change the intrinsic properties of other objects. We never thought this about singular reference. Referring to

something does not bring it into existence. Why should managing to include it in a domain of quantification?

So, I do not think my position leads to anti-realism. Even so, my position does, I think, remind us of some important features of realism as a philosophical doctrine. I have in mind particularly views that take realism as a *semantic* thesis, as I believe does Putnam, and certainly as do Michael Dummett and his followers. The idea, roughly again, is that realism is a thesis about the nature of our claims. Our claims are, to such a realist, objective in some sense. They are objectively about the world, and objectively right or wrong. This entails that somehow or another the world we talk about is equally objective. Call realism, formulated as a semantic thesis, *semantic realism*. (This is realism about the world, formulated in semantic terms, not realism about the semantic.)

I am going to examine this way of formulating realism, not defend it. But I do want to mention one example due to Crispin Wright [1992], that helps to show why it might be an attractive approach, or at least, more attractive than formulating realism as a specifically ontological thesis. We are somewhat inclined, though not entirely so, to be anti-realist about the comic. We think that to a great extent, what is funny is dependent on the agent finding it funny. But not quite entirely. We sometimes feel that there must be at least some cases where things are just plain not funny, and to find them so is to make an error. But there is no question of ontological commitment here. No one ever suggested the issue comes down to whether there are ‘funnies’—objects that are ‘funny-makers’.

But if realism is not a matter of ontology, what is it? ‘Objective’ and ‘factual’ are nice metaphors, but do not get us very far. Two proposals have been hazarded, both semantic. From

Putnam, we have the idea that substantial notions of reference and truth can be appealed to to explain the important features of objective claims.<sup>19</sup> From Dummett, we have an even more direct proposal that realism is a matter of semantics. A typical example (among many) is:

We may thus characterize a realistic interpretation of a given class of statements as one which applies to them, in accordance with the structure they appear on the surface to have, the classical two-valued semantics, in particular treating the apparent singular terms in them as denoting objects (in the relevant domain) and the statements themselves as being determinately true or false [1991b, p. 246].

Against semantic realism goes, of course, semantic anti-realism, which normally involves some features of epistemic, or verificationist, or theory relative views of truth or reference.

I do not want to defend any of these proposals here, only to examine whether my claims about quantification, and more generally about interpretation, have any impact on semantic realism. One might worry that the quote from Dummett already is in tension with my view. I said that interpretation is risky, and that some of the risk derives from the possibility of finding ourselves in paradoxical discourse situations. But if so, it seems, then the contents of our claims is in part a matter of complex features of utterance and interpretation. But these are surely in part features of us, as speakers and as agents. Is this not precisely the kind of dependence on agents that anti-realism postulates, which undermines determinate truth value?

No. Not at all. Consider first a more familiar case of a demonstrative phrase. I say ‘that table has a paper on it’. What I say is determined substantially by how we interpret ‘that table’ on the occasion of utterance. This is normally a matter of which table we take to be salient in a given conversational context (and perhaps a number of other factors as well, such as speaker’s intentions). Hence, what we talk about by saying ‘that table’ is to a great extent determined by features specific to agents on occasions of utterance. But there is hardly a threat of anti-realism

here. The fact that what we talk about in this case is in part a matter of what agents find salient does not imply that the reference relation is non-objective, nor does it imply that the truth of the claim we make is an agent-relative or epistemic matter. It does not lead us to deny the objectivity of the world! That would be absurd.

The quantification case is basically the same. Quantification, I have suggested, works a bit more like singular reference than we might have thought. It does so in that what domain we wind up quantifying over is in some ways affected by aspects of interpretation, including linguistic features of a discourse. But this does not give the meanings expressed by using quantifiers any kind of anti-realist features. Once interpretation is carried out, once it is fixed what is said, truth conditions and truth and reference may be perfectly objective.

The proposal that we are led to anti-realism is confused, but the confusion is invited by the way semantic realist positions are often put. If we understand Dummett's call for a particular kind of semantics for statements as requiring that the connection between sentences, or more properly utterances, and their contents be determined independently of speakers, we get a claim that conflicts with my view on quantification. But we also get a view that is outright silly. *Of course*, what is said in a particular utterance is a matter of language and speaker. After all, it is a speaker, speaking a language, who says something in the first place.

I doubt anyone ever proposed this rather absurd view. But to make it more clear what semantic realism is supposed to be, a distinction will be helpful. We must distinguish what I shall call *semantic interpretation* from *semantic foundations*. Questions of semantic interpretation are those of the relation between utterance and meaning: what sorts of propositions are expressed by what sorts of utterances, in which contexts, and why. Typical questions of

semantic interpretation include those about the semantic properties of particular sorts of terms, like names, or quantifiers, and those about the way context determines what is said by particular utterances, such as those involving indexicals and demonstratives. Questions of semantic foundations are those of the underlying nature of the basic semantic notions, such as proposition, meaning, reference, and truth. Typical questions of semantic foundations include that of whether there is a substantial or objective relation or reference, and whether truth conditions attaching to utterances are epistemic in nature.

The questions I mentioned as typical of semantic foundations are those that are at issue for semantic realism. There is no question about realism or anti-realism at the level of semantic interpretation. We should expect semantic interpretation to be in part determined by the speaker—it is the speaker doing the speaking! This should be obvious. It might be something we can forget when concentrating on sentences, which are abstract objects. It is utterances, however, that express propositions; only in some degenerate cases can we think of sentences as doing this job. Utterances are actions taken by speakers, so we had better expect a contribution of the speaker to the relation between utterance and meaning. But that speakers have a say in what they say does not make semantic *foundations* determined by what they can say or do. Speakers are involved in expressing propositions, in making meaningful claims, but we may not conclude from this that the claims made have only speaker-relative or epistemic truth conditions. To do so is to confuse the proposition with the act of expressing the proposition. Nor can one conclude that the relation between propositions the world is speaker-dependent. To do so would be to confuse the relation between the proposition and the world with the act of expressing the proposition.

Semantic foundations and semantic interpretation are to a great extent autonomous, and semantic realism must be construed as a thesis about semantic foundations. Once we see this distinction, I think we can see a little better why even the rather drastic conclusion about quantification I have been defending does not threaten realism. It is a thesis about semantic interpretation, about which particular kinds of claims can be made by utterances involving unrestricted quantifiers. It is not a claim about semantic foundations at all.

Nonetheless, I do think that the thesis I have advanced is a bit surprising, though I hope I have rendered it more plausible than it might have seemed at the outset. Because of the extensibility of our maximally broad conception of object, for any fundamental domain used to interpret an unrestricted quantifier, there is a strictly broader domain. Hence, no utterance can accomplish absolutely unrestricted quantification. I have suggested this is not so implausible, as it is only in very special cases that we might even want to try, and these may be given special treatment. I have also argued that this result does not have large-scale metaphysical implications.

What is the moral, if not anything for realism? Interpretation is an activity—it is something we do. We have uncovered some limits to how it can be carried out. They reveal not instability in the world of objects, but an instability in our activity of interpreting. Sometimes, typically in situations where discourse becomes paradoxical, we have to revise our scheme of interpreting, by expanding the background domain we use to interpret. In this sense, the limitations built into interpretation make it inherently *risky*. Anyone who has ever tried to communicate knows that it is risky in lots of ways, but I have drawn attention to one we might not have considered. It has a surprising upshot: no domain of quantification can be absolutely everything.

## Notes

1. I am here suppressing issues of context dependence, as well as those of non-assertoric utterances.
2. I have investigated some aspects of this in my [MS].
3. I have borrowed the example of ‘heavy’ from Tappenden [1993]. A similar point is stressed in the classic discussion of Wright [1975].
4. Notably Frege [1884], among many of his writings. My own discussion owes a great deal to Parsons [1982]. Views of Quine, e.g. Quine [1969], have some affinity to the one I sketch, though as quantification is precisely what is at issue here, the more Fregean approach applies more directly. One sees hints of logical characterizations of the notion of object long before Frege. For instance, Aristotle says in *Categories* (3b10), “Every substance seems to signify a certain ‘this’ ” [Barnes, 1984, p. 6]. I would not, though, go quite so far as to really attribute this very modern idea to ancient writers.
5. For some discussion of these matters, see Dummett [1981], Wright [1983], and Hale [1987].
6. Williamson’s version of Russell’s paradox is the most general with which I am familiar. The observation that one does not need sets *per se* may also be found in Parsons [1974], which uses arguments based on the Liar paradox to reach similar conclusions for propositions, and for the notion of truth relative to a domain of quantification. For a detailed comparison of the Liar and Russell’s paradox reasoning, see Feferman [1984].
7. The original description of the plural interpretation in e.g. Boolos [1984] relied crucially on translation to natural-language plural idioms.
8. Though I suspect that many find the thesis I advocate implausible, one of the few to argue against it explicitly is Cartwright [1994]. There is much that I agree with in his discussion, especially his point that Russell’s paradox cannot by itself sustain my conclusion. Indeed much of my discussion is tailored to try to answer his. He identifies what he calls the ‘all-in-one’ principle, which holds that if it is coherent to quantify over some objects, there must be a single object which contains them. A substantial portion of his paper is directed towards rejecting this principle.

In the class-based version of Russell’s paradox, I do collect all into one. But as I have tried to make clear, I am not relying on the all-in-one principle *per se*. The issue as I see it is not one of collecting all into one, but of nominalization. Indeed, in the more general version of Russell’s paradox, there is no need to collect all into one in the sense of forming a class. Rather, what is needed is to nominalize the results of the process of interpretation. I have tried to argue here that we have good reason to accept such nominalization. The process of interpretation itself shows us what the referent of such a term is, and the widespread presence of natural-language nominalization underwrites our confidence that we do indeed have a genuine singular term. I

thus propose that to the extent I do wind up collecting all into one, it is for good reasons, not based upon a general all-in-one principle.

It may be that the logical notion of object, together with widespread nominalization and the constraints on the interpretation of quantifiers, imply a modest form of an all-in-one principle. If one succeeds in quantifying over some objects, one can nominalize the domain, which provides an object which we may (at least loosely) say collects all into one. If this is so, it is primarily because of the high demands on the interpretation of quantifiers. Providing a determinate specification of a domain is enough to license successful nominalization. But this is a feature of the process of interpretation as much as a principle about classes or collections.

9. Williamson [MS] himself takes a slightly different line. He proposes that interpretation is done in a second-order language that is, as we might say, *primitively* second order. It crucially is not a language we can learn by translation into our own natural language, for that would license just the sort of nominalization that must be avoided. Hence, for instance, the plurals interpretation of Boolos [1984] is not available. Nor can this language be interpreted in the usual class-theoretic ways. Hence, it is a direct cognitive achievement to grasp this sort of language.

I appreciate that Williamson meets the challenges of defending absolutely unrestricted quantification head-on. But I still maintain the costs of this route are too great. I do not see how this language could ever be learned, as it cannot be learned by our usual routes to acquiring language, or by way of some sophisticated mathematics. I especially do not see how it could be deployed by speakers in interpretation, without embedding it in a framework which would recover nominalization. As a native speaker and as one who grasps some class theory, I can take this language and nominalize its predicate positions; correctly, as far as I can tell. To insist I cannot is not only to make this language difficult to learn, but to make it oddly isolated from the rest of our cognitive architecture. Williamson's approach is an elegant and direct way to make space for absolutely unrestricted quantification, but again, one which leads to a radical position indeed. Too radical, I think, and radical enough to make the restricted quantifiers options I offer in Section (II) appear all the more plausible.

10. I should caution against inferring from the indefinite extensibility of the logical notion of object that it is vague. I did indeed describe the phenomenon of extensibility as related to vagueness, as both vague and indefinitely extensible concepts display the same kind of open-endedness. But this feature may well not be sufficient for vagueness.

11. The iterative conception of set (e.g. Boolos [1971]) may be taken as an attempt to provide a determinate specification of the domain of sets. It certainly has the form of an inductive specification, as I suggested we might expect of specifications of fundamental domains. The well-known difficulties surrounding the iterative conception (as in Boolos [1971, 1989]) make it doubtful that the task of making sense of quantification over sets has been completed, but I see no reason to hold it cannot be done.

12. Syntactically, quantifiers like 'every' are determiners, which combine with a common noun (or an N') to form a noun phrase. The common noun restricts the quantifier. Semantically, these determiners are interpreted as binary generalized quantifiers, making NPs restricted monadic

quantifiers. Following Barwise and Cooper [1981], it is common to treat NPs like ‘everything’ as having the same structure, at least at the appropriate level of logical form. They see ‘everything’ as having the structure  $[[\text{every}]_{\text{DET}}[\text{thing}]_{\text{N}}]_{\text{NP}}$ . For a more recent survey, and some complications to the picture, see Keenan and Westerståhl [1997].

13. For discussion of contextual restrictions of quantifier domains, see Westerståhl [1985] or von Fintel [1994] or Stanley and Szabó [2000].

14. I pursue some aspects of this in my [MS].

15. My own suggestions here have some affinity to those of Russell [1908] and Parsons [1974]. Russell [1908, pp. 96–97] suggests that talk of cardinality can be typically ambiguous because there is a 1-1 correspondence between classes of corresponding cardinality but different type. My own position has a further affinity with Russell’s. He writes [1908, p. 73], “[If a propositional function] ceases to be significant when the variable goes outside of a certain range, then the variable is *ipso facto* confined to that range, without the need for any explicit restriction to that effect.”

16. One other extraordinary measure, advocated by Michael Dummett [1981, 1991a, 1993], is to treat some unrestricted quantifiers as constructive. I do not favor this approach, but it is worth noting that it is open even to those who are not convinced by Dummett’s general arguments against classical logic in favor of intuitionistic logic. For a discussion of Dummett’s view targeted specifically at set theory, but sympathetic to the position I argue here, see Clark [1993].

17. A version of this objection is voiced by McGee [2000]. He worries that without absolutely unrestricted quantification, we have reason to doubt that “a science such as ontology is even possible” [p. 54]. Objections of a similar kind were raised by an anonymous referee.

18. For instance, Rosen [1993] says so explicitly. Dummett argues in favor of anti-realism by way of indefinite extensibility, which reveals a significant difference in the ways we approach this notion.

19. The idea that the reference relation is key to realism appears in many of Putnam’s writings. Putnam [1978a] makes it explicit (building on ideas of Richard Boyd).

## References

Barnes, J. (ed.), 1984. *The Complete Works of Aristotle*, Princeton, Princeton University Press.

Barwise, J. and R. Cooper, 1981. “Generalized Quantifiers and Natural Language,” *Linguistics and Philosophy* 4: 159–219.

- Boolos, G., 1971. "The Iterative Conception of Set," *Journal of Philosophy* **68**: 215–232. Reprinted in Boolos (1998).
- , 1984. "To Be is to Be a Value of a Variable (or to Be Some Values of Some Variables)," *Journal of Philosophy* **81**: 340–449. Reprinted in Boolos (1998).
- , 1989. "Iteration Again," *Philosophical Topics* **42**: 5–21. Reprinted in Boolos (1998).
- , 1998. *Logic, Logic, and Logic*, Cambridge, Harvard University Press.
- Cartwright, R., 1994. "Speaking of Everything," *Nous* **28**: 1–20.
- Clark, P., 1993. "Sets and Indefinitely Extensible Concepts," *Aristotelian Society Supplementary Volume* **67**: 235–249.
- Davidson, D., 1973. "Radical Interpretation," *Dialectica* **27**: 313–328. Reprinted in Davidson (1984).
- , 1984. *Inquiries into Truth and Interpretation*, Oxford, Oxford University Press.
- Dummett, M., 1981. *Frege: Philosophy of Language*, second edition, Cambridge, Harvard University Press.
- , 1991a. *Frege: Philosophy of Mathematics*, Cambridge, Harvard University Press.
- , 1991b. *The Logical Basis of Metaphysics*, Cambridge, Harvard University Press.
- , 1993. "What Is Mathematics About?" in *The Seas of Language*, pp. 429–445, Oxford, Oxford University Press.
- Feferman, S., 1984. "Towards Useful Type-Free Theories, I," *Journal of Symbolic Logic* **49**: 75–111. Reprinted in Martin (1984).
- von Fintel, K., 1994. *Restrictions on Quantifier Domains*, Ph.D. Thesis, University of Massachusetts at Amherst.
- Frege, G., 1884. *Die Grundlagen der Arithmetik*, Breslau, Verlag von Wilhelm Koebner. English translation by J. L. Austin as *The Foundations of Arithmetic*, Evanston, Northwestern University Press, 1980.
- Glanzberg, M., MS. "A Contextual-Hierarchical Approach to Truth and the Liar Paradox," manuscript.

- Hale, B., 1987. *Abstract Objects*, Oxford, Basil Blackwell.
- Keenan, E. L. and D. Westerståhl, 1997. “Generalized Quantifiers in Linguistics and Logic,” in J. van Benthem and A. ter Meulen (eds.), *Handbook of Logic and Language*, pp. 837–893, Cambridge, MIT Press.
- Kripke, S., 1975. “Outline of a Theory of Truth,” *Journal of Philosophy* **72**: 690–716. Reprinted in Martin (1984).
- Lewis, D., 1975. “Languages and Language,” in K. Gunderson (ed.), *Minnesota Studies in the Philosophy of Science*, volume VII, pp. 3–35, Minneapolis, University of Minnesota Press. Reprinted in Lewis (1983).
- , 1983. *Philosophical Papers*, volume I, Oxford, Oxford University Press.
- Martin, R. L., 1984. *Recent Essays on Truth and the Liar Paradox*, Oxford, Oxford University Press.
- McGee, V., 2000. “‘Everything’,” in G. Sher and R. Tieszen (eds.), *Between Logic and Intuition*, pp. 54–78, Cambridge, Cambridge University Press.
- Parsons, C., 1974. “The Liar Paradox,” *Journal of Philosophical Logic* **3**: 381–412. Reprinted in Martin (1984).
- , 1982. “Objects and Logic,” *The Monist* **65**: 491–516.
- Putnam, H., 1978a. *Meaning and the Moral Sciences*, London, Routledge and Kegan Paul.
- , 1978b. “Realism and Reason,” in *Meaning and the Moral Sciences*, pp. 123–140, London, Routledge and Kegan Paul.
- , 1981. *Reason, Truth and History*, Cambridge, Cambridge University Press.
- Quine, W. V. O., 1969. “Existence and Quantification,” in *Ontological Relativity and Other Essays*, pp. 91–113, New York, Columbia University Press.
- Rosen, G., 1993. “The Refutation of Nominalism (?),” *Philosophical Topics* **21**: 149–186.
- Russell, B., 1908. “Mathematical Logic as Based on the Theory of Types,” *American Journal of Mathematics* **30**: 222–262. Reprinted in Russell (1956).
- , 1956. *Logic and Knowledge*, R. C. Marsh (ed), London, George Allen and Unwin.

- Stanley J. and Z. G. Szabó, 2000. "On Quantifier Domain Restriction," *Mind and Language* **15**: 219–261.
- Tappenden, J., 1993. "The Liar and Sorites Paradoxes: Towards a Unified Treatment," *Journal of Philosophy* **90**: 551–577.
- Westerståhl, D., 1985. "Determiners and Context Sets," in J. van Benthem and A. ter Meulen (eds.), *Generalized Quantifiers and Natural Language*, pp. 45–71, Dordrecht, Foris Publications.
- Williamson, T., MS. "Everything," manuscript.
- Wright, C., 1975. "On the Coherence of Vague Predicates," *Synthese* **30**: 325–365.
- , 1983. *Frege's Conception of Numbers as Objects*, Aberdeen, Aberdeen University Press.
- , 1992. *Truth and Objectivity*, Cambridge, Harvard University Press.