

QUANTIFIER VARIANCE AND REALISM

Eli Hirsch
Brandeis University

A pervasive theme in Hilary Putnam's writings for many years, running as a constant thread through various changes in his views about realism, is a doctrine that he calls "conceptual relativism", representative formulations of which are as follows:

[A]ll situations have many different correct descriptions, and ... even descriptions that, taken holistically, convey the same information may differ in what they take to be 'objects' ... [T]here are many usable extensions of the notion of an object ...¹
[T]he logical primitives themselves, and in particular the notions of object and existence, have a multitude of different uses rather than one absolute 'meaning'.²

In these passages Putnam seems to be saying that the quantificational apparatus in our language and thought—such expressions as "thing", "object", "something", "(there) exists"—has a certain variability or plasticity. There is no necessity to use these expressions in one way rather than various other ways, for the world can be correctly described using a variety of concepts of "the existence of something". One of his favorite examples concerns a disagreement between mereologists and anti-mereologists as to how many objects there are in some domain.³ Suppose we are evaluating the truth of the sentence, "There exists something that is composed of Clinton's nose and the Eiffel Tower". Mereologists will accept this sentence, whereas anti-mereologists will reject it. Putnam's doctrine of quantifier variance implies that the expression "there exists something" can be interpreted in a way that makes the sentence true or in a way that makes the sentence false. Since both interpretations are available to us, we have a choice between operating with a concept of "the existence of something" that satisfies the mereologist or operating with a different concept that satisfies the anti-mereologist.

The doctrine of quantifier variance may be philosophically unsettling. Our initial reaction may be that, if we are free to choose between different ways of conceiving of "the existence of something", then this threatens a robust realist

sense that there are things in the world whose existence does not in any way depend on our language or thought. The nature and force of this threat is one of the topics of this paper.

Before proceeding let me emphasize that this paper is not an exercise in Putnam-exegesis. I take it as obvious that the doctrine of quantifier variance is a central part of Putnam's overall position, but I make no attempt here to trace Putnam's evolving views as to how this doctrine relates to different versions of realism. The possible threat to realism posed by this doctrine is something that I am trying to work through on my own terms, though I have little doubt that virtually every point I am going to make can be found somewhere in Putnam's writings.

I

I think we should begin by making sure to repudiate a thoroughly confused, though somehow tempting, formulation of what the threat to realism is: "Since, according to the doctrine of quantifier variance, our linguistic decisions determine whether or not there exists something composed of Clinton's nose and the Eiffel Tower, it evidently follows that this thing's existence or non-existence depends on our language and thought. If this point generalizes to every application of our concept of 'the existence of a thing'—as the doctrine seems to imply—then quantifier variance evidently conflicts with the realist idea of things existing independently of language and thought".

The fallacy in this formulation lies in the claim that the doctrine of quantifier variance implies that our linguistic decisions determine whether or not there exists something composed of Clinton's nose and the Eiffel Tower. What the doctrine does imply is that our linguistic decisions determine the meaning of the expression "there exists something"; hence, they determine the meaning of the sentence "There exists something composed of Clinton's nose and the Eiffel Tower". Hence, the truth or falsity of this sentence depends in part on our linguistic decisions. It is merely a use-mention confusion to conclude that whether or not there exists something composed of Clinton's nose and the Eiffel Tower depends on our linguistic decisions.

Consider the two sentences: "There exists something composed of Clinton's nose and the Eiffel Tower" and "Whether or not there exists something composed of Clinton's nose and the Eiffel Tower depends on our linguistic decisions". Quantifier variance implies that the expression "there exists something" can be interpreted in a way that makes the first sentence true or in a way that makes it false. But there is no relevant way to interpret "there exists something" that would make the second sentence true. The second sentence expresses an absurd form of linguistic idealism that is not at all implied by quantifier variance.

There is the familiar joke about how many tails a dog would have if the word "tail" were used to refer to legs—the correct answer, everyone seems to

agree, is “one”, because how many tails a dog has does not depend on our linguistic decisions. The meaning and truth-value of the sentence “A dog has one tail” depends on our linguistic decisions, but how many tails a dog has does not depend on our linguistic decisions. Why does it seem harder to grasp this point when we are talking about the meaning of quantifier expressions rather than a general term such as “tail”? I think that part of the difficulty stems from our wanting to say, when we are formulating the doctrine of quantifier variance, that the relevant variations would still leave us with a *kind of quantifier expression*, an expression that, as Putnam put it in the earlier quote, continues to signify *a notion of the existence of something*. This formulation may be misconstrued as implying that, although the meaning of the quantifier expressions and the notion of existence remain fixed, what we are somehow going to accomplish with our linguistic decision is to alter the truth-values of our existential claims. And that might indeed amount to a lunatic form of linguistic idealism. What needs to be stressed, however, is that the doctrine of quantifier variance only allows for the possibility of a *change in the meaning* of quantifier expressions, yielding a *different* or *extended* notion of existence; only in this way does the doctrine allow our linguistic decisions to affect the truth-values of existential sentences. Once the meaning of the quantifier is fixed there is no further effect that our decisions can have on the truth-values of typical existential sentences.

To say that even after there has been a change in the meaning of quantifier expressions we still have a “kind of quantifier expression” and a “notion of existence” is merely to indicate a degree of similarity between the concepts we started with and those we end up with. Nothing is being said here to imply the idealist view that what exists in the world depends on our linguistic or conceptual decisions. Suppose that I start out with the anti-mereologist’s position that the sentence “There exists something composed of Clinton’s nose and the Eiffel Tower” is false. If I accept quantifier variance I will allow that I can make intelligible to myself a change in the meaning of the expression “there exists something” which would have the effect of rendering the sentence true. I might characterize this change as simply “giving up the quantifier” and “giving up the notion of the existence of something”, but it seems more natural to characterize it rather as “acquiring a new kind of quantifier” and “acquiring a new notion of the existence of something”. The second characterization seems more natural because it seems clear that the imagined change in the meaning of the expression “there exists something” will leave the expression’s general role in the language largely intact. In particular, the purely syntactic and formal logical properties of the expression will not be changed at all (the formal principles of quantificational logic will be unaltered). It therefore seems natural to follow Putnam in treating relevant variations in the meaning of such expressions as “there exists something” as yielding an altered quantificational apparatus and an altered concept of the existence of something.

If I start out with the anti-mereologist’s stance, what exactly is involved in changing the meaning of the quantifier with the effect of making the mereolo-

gist's sentences come out true? There are important complications in the answer to this question that I will return to later, but the basic idea is quite simple. In general, we explain the meaning of a logical constant by describing the role it plays in determining the truth-conditions of sentences. Thus we explain the meaning of "and" by saying that sentences of the form " p and q " are true if and only if both the sentence " p " and the sentence " q " are true. If we were to explain some imagined change in the meaning of "and" we would do so by describing a change in the truth-conditions of sentences containing "and". Analogously we explain the relevant change in the meaning of the quantifier, which will render the mereologist's sentences true, roughly as follows: In the new meaning, any sentence of the form "There exists something composed of the F -thing and the G -thing" is true if the expression "the F -thing" refers to something and the expression "the G -thing" refers to something. We are here using the quantifier in one of its meanings to explain the other meaning. It is evidently not part of the doctrine of quantifier variance to claim that the meaning of the quantifier can be somehow analyzed or defined in terms not involving the quantifier—any more than the meaning of "and" can be analyzed or defined in terms not involving conjunction. The issue of quantifier variance should therefore not be conflated with familiar questions about the analyzability of (the criteria for) the identity or existence of a thing. Quantifier variance is not a matter of substituting one "definition" for another; it's a matter of substituting one range of truth-conditions for another.

II

I suspect that the deepest source of the illusion—and I am maintaining that it is an illusion—that quantifier variance conflicts with realism stems from the analogy we are led to draw between different kinds of quantifiers, the analogy, that is, that leads us to speak, not simply about eliminating the quantifier and the concept of existence, but about a variation that still leaves us with a different quantifier and a different concept of existence. Suppose, again, that I start out with the anti-mereologist's position that the sentence "There exists something composed of Clinton's nose and the Eiffel Tower" is false. In my use of "there exists something" a sentence of the form "There exists something composed of the F -thing and the G -thing" counts as true only if "the F -thing" and "the G -thing" refer to things that are connected (united) in some special ways. Let me call my use of "there exists something" the A-use. I now imagine a different use of "there exists something"—I will call this the M-use—in which a sentence of that form counts as true so long as "the F -thing" and "the G -thing" refer to things, no matter how they are connected. Now, how am I to describe in my language, the A-language, what is going on in the M-language?

When speakers of the M-language assert "There exists something composed of Clinton's nose and the Eiffel Tower" they speak the truth. That is implied by the semantic rule that has been explicitly laid down in my description

of the truth-conditions in the M-language of sentences of the form “There exists something composed of the *F*-thing and the *G*-thing”. I assume, however, that this explicit stipulation carries with it a natural way of filling in the truth-conditions for an indefinite variety of other sentences in the M-language. Take, for example, the sentence “There exists something that is now being touched by exactly two people”. This sentence is not rendered true in my A-language by a situation in which one person is touching Clinton’s nose and one person is touching the Eiffel Tower,⁴ but my implicit assumption is that the sentence would be true in the M-language with respect to that situation. When I reflect on a range of examples of this sort I find myself tempted to say (in my A-language) something like the following: Although there doesn’t exist anything composed of Clinton’s nose and the Eiffel Tower, it’s somehow *as if* there does exist such a thing relative to the M-language. This “as if” formulation is a simile; it is a way of expressing the felt holistic analogy between the A-use of the quantifier expressions and the M-use. As such it may be perfectly innocuous. But similes can often lead to metaphors, and metaphors, if we are not careful, can sometimes lead to lunacy. The move from simile to metaphor might take the form of simply dropping the words “as if” from the previous formulation. We then wind up saying, “Relative to the M-language, there exists something composed of Clinton’s nose and the Eiffel Tower”. As a metaphor this may still be okay; it may in fact qualify as what John Wisdom called an “illuminating paradox”.⁵ It can, however, turn into utter confusion and philosophical madness if one forgets that it is merely a metaphor. One may then wind up claiming that language or thought literally creates everything that exists, that nothing could exist if there were no people speaking or thinking. What we need to be clear about—and this may require a continual effort on our parts—is that the doctrine of quantifier variance does not imply any such idealist formulation.⁶

The push towards the simile and metaphor, and the accompanying risk of falling off the edge of idealist madness, is enhanced when we consider the following kind of sentence: “In the M-language the expression ‘thing composed of Clinton’s nose and the Eiffel Tower’ refers to a thing composed of Clinton’s nose and the Eiffel Tower”. This sentence (about the M-language) cannot be true in my A-language. Since (as I say in my A-language) there isn’t any thing composed of Clinton’s nose and the Eiffel Tower, there can’t be any such thing referred to by any expression in the M-language. But if we make the natural assumption that the truth-conditions for sentences in the M-language containing the word “refers” will respect the disquotational principle, then the previous sentence about the M-language, although false in the A-language, counts as true in the M-language. To take this one small step further, the sentence “In the M-language the expression ‘thing composed of Clinton’s nose and the Eiffel Tower’ refers to something” counts as false in the A-language but true in the M-language.

As goes “existence” so goes “reference”. The lesson here is that if we are imagining that the quantifier expressions in the M-language function differ-

ently from our A-quantifiers, then we can—and naturally will—imagine a correlative difference in the use of the word “refers”. If we alter our concept of “a thing” then we alter our concept of “reference to a thing”. But here especially we may feel impelled to say, “It’s as if, relative to the M-language, there is something being referred to by the expression ‘thing composed of Clinton’s nose and the Eiffel Tower’”. Again we must struggle to keep this simile in tow and not let it lead us into idealism.

III

Let me now briefly consider a few formulations that figure in the literature on realism and ask how they relate to quantifier variance. I will assume that these formulations are being posed in the A-language. (In the next section I will argue that people who speak plain English are in fact using the A-language, not the M-language.)

1. *The truth of a (contingent) statement in any language depends on what things exist and what properties these things have.* This claim is not threatened by quantifier variance. People who speak the M-language will have a different concept of “a thing” than mine, but the truth of their statements will depend on how it is with things in the world, in my (A-language) sense of “things in the world”. For instance, if one person is touching Clinton’s nose and another person is touching the Eiffel Tower, then the existence and properties of these things will render true the M-statement, “There exists a thing that is being touched by exactly two people”.

Might there, however, be truths that do not depend on how things are, truths that are somehow not about things in my sense (or perhaps in any sense that I can make intelligible to myself)? I think that this question takes us into the vicinity of Kantian noumena and mysticism. It’s not a question that I’m addressing here.

This last question should, however, be distinguished from another one: Might there be things whose nature and unity I cannot understand? The doctrine of quantifier variance does not preclude there being such things.⁷ (This relates to my earlier point that the doctrine is not to be confused with issues about the “analyzability” of existence and identity.)

2. *The truth of any (contingent) statement in any language depends on the existence and properties of the things referred to by the non-logical expressions in the statement.* This claim as it stands is immediately refuted by such sentences as “The average professor has fewer children than the average plumber”, the truth of which evidently does not depend on there being any things referred to by the expressions “the average professor” and “the average plumber”. Let’s suppose, however, that this sentence can be viewed as in some sense “merely a transformation” of some other “basic” sentence. Claim 2 might then be roughly understood as saying that, in any language, the truth of a basic sentence depends on the existence and properties of the things referred to by the

non-logical expressions, and the truth of other sentences depends on their equivalence by way of transformation rules to some basic sentences.

Is claim 2 threatened by quantifier variance? If we suppose that the notion of “reference to a thing” is kept fixed to the language in which 2 is formulated (which I am now imagining to be the A-language), then I suspect that 2 is threatened. In terms of my A-concept of “reference to a thing” I cannot explain the truth of the M-statement “There exists a thing that exactly two people are touching” by appealing to the reference of “touching”.⁸ (Quantifier variance might thus be said to induce a certain kind of systematic difference of meaning in the word “touching” and, by the same token, virtually any other general word.⁹) I might try to treat as “basic” only those sentences of the M-language that have the same truth-conditions as homophonic sentences in the A-language and derive the rest by transformations, but it is not obvious what such transformation rules would look like. Claim 2 requires that it be possible to formulate in my A-language something approximating to a Tarski-style theory of truth for the M-language, that is, very roughly put, a finitary account in terms of reference relations that yields for each of the indefinite number of M-sentences what the conditions are for its truth. It seems quite possible that this requirement can’t be met.

If claim 2 constitutes a certain kind of “correspondence theory of truth” (what might be called a “referential correspondence theory”) then this kind of correspondence theory may indeed fall to quantifier variance—assuming, again, that claim 2 is being formulated in the A-language and that what is meant by “reference to a thing” is kept fixed to this language. But I think it is therefore clear at this point that this kind of correspondence theory is not essential to a straightforward realist view of the world. The possible falsity of claim 2 does not threaten the basic realist idea (now being expressed in the A-language) that the world consists of things whose existence and properties are independent of language or consciousness.

Putnam takes Donald Davidson’s animadversions against “the very idea of a conceptual scheme” to constitute a repudiation of quantifier variance.¹⁰ I don’t fully understand what Davidson is driving at in his talk against “conceptual schemes”, but I doubt that it has really any definite connection to the issue of quantifier variance. Let us note that the M-language is “translatable” into my A-language in at least the following sense: For any M-sentence I can find an A-sentence with the same truth-conditions, where two sentences have the same truth-conditions if, relative to any context of utterance, they hold true with respect to the same possible situations.¹¹ This kind of intertranslatability between the A-language and the M-language holds even if it’s impossible to formulate in the A-language a finitary theory of truth for the M-language. (“For any M-sentence I can find an A-sentence with the same truth conditions” does not entail “There is a finitary truth-theory which yields, for any M-sentence, an A-sentence with the same truth-conditions”.) It is indeed because of this intertranslatability that the proponent of quantifier variance maintains that the two languages are equally capable of truthfully describing the world. I don’t know,

therefore, whether the A-language and the M-language qualify as “different conceptual schemes” in Davidson’s sense. If they do, he has certainly said nothing to show why there couldn’t be different conceptual schemes in this way.

Another idea associated with Davidson is that any (learnable) language must be describable in terms of (something approximating to) a Tarski-style truth-theory.¹² Interpreted in one way, this claim, too, is not threatened by quantifier variance. If it is possible to formulate in the A-language a truth-theory for the A-language, then it is possible to formulate in the M-language a truth-theory for the M-language, each theory being formulated in terms of each language’s meaning of “reference to a thing”. If, however, Davidson is claiming that it must be possible to formulate in our (“home”) language a truth-theory for any possible language, this seems clearly untenable, for we can surely conceive of people whose sensory apparatus differs from ours to the extent that we cannot describe the truth-conditions of some of their sentences.¹³

A question that might be raised, however, is this: If we start out speaking the A-language (as I think we in fact do), and we are not able to formulate in this language a truth-theory for the M-language, how can we come to learn that language? How do we come to agree (as I assume we do) on the truth-conditions of an indefinite number of M-sentences? The answer is that we are relying on our shared sense of the analogy between our A-quantifier and the M-quantifier. Starting with the A-language I teach someone the M-language *ostensively*, by giving a few representative examples of how the M-language works. That anyone can then be expected to go on in the same way is part of what Putnam means when he says that our quantificational concepts admit of “many usable extensions”.¹⁴

3. *The truth of any (contingent) statement in any language depends on what the facts are in the world.* This might also be put by saying that a statement’s truth depends on *the way the world is*, the statement’s being true only if the world is the way the statement says the world is. Claim 3 is another perennial version of “the correspondence theory of truth”. Is it threatened by quantifier variance?

In the philosophical literature we find facts (as well as states-of-affairs, propositions, and properties) sometimes understood in a coarse-grained *unstructured* manner and sometimes in a fine-grained *structured* manner.¹⁵ A structured fact is what Putnam describes—and rejects—as a “sentence-shaped thing in the world”.¹⁶ A fact in this sense is built up in a certain way out of things and properties (and perhaps logical operations). If we have logically equivalent sentences, such as “This is round” and “This is either round and red or round and not red”, then each sentence expresses a different structured fact, or, as another alternative, they express the same structured fact but only one sentence (the first) succeeds (or succeeds better) in properly picturing the fact’s structure. In the coarse-grained sense, however, these sentences express the same unstructured fact, and it therefore makes no sense to ask which sentence does better at depicting this fact.

If claim 3 is understood in terms of unstructured facts it does not conflict at all with quantifier variance. Quite the contrary, the basic idea of quantifier variance can be nicely formulated by saying that the same (unstructured) facts can be expressed using different concepts of “the existence of a thing”, that statements involving different kinds of quantifiers can be equally true by virtue of the same (unstructured) facts in the world.

The notion of a structured fact does, however, raise certain problems for quantifier variance—but one must be careful not to misunderstand what these problems are. If I change what I mean by “a thing” then I must also change what I mean by “the way a fact is built up out of things and properties”; hence, I must change what I mean by “a structured fact”. But, note carefully, that it does not follow that there are no structured facts independent of language—any more than it follows that there are no things independent of language. Indeed, when one looks at this carefully I think one sees that claim 3 can be sustained even if the facts on which the truth of statements depend are taken to be structured. I can say in my A-language that the truth of an M-statement depends on the structured facts, even though what I mean by “the structure of the facts” is not what a speaker of the M-language means.

A problem arises, however, if one wants to say that corresponding to any true statement in any language there is, as Putnam puts it, “a *unique* sentence-shaped thing in the world”,¹⁷ in other words, that each true sentence states one structured fact. How can we say which one that is? If we stick to our own language then we have the trivial disquotational formulation, “The sentence ‘*p*’ states the structured fact that *p*”, and perhaps we don’t have to worry further about which structured facts are identical with which. But if I am speaking the A-language, how am I to say which structured fact is stated by the true M-sentence “There exists a thing that exactly two people are touching”, when I am able to “translate” that sentence into a variety of structurally different sentences in my language having the same truth-conditions? I don’t doubt that some moves might be made here (for instance, it might be suggested that it is indeterminate which structured fact, in my A-sense of “structured fact”, the true M-sentence states), but I am inclined to agree with Putnam that, once we’ve accepted quantifier variance, there is no point in trying to hold onto language-shaped facts that are in the world independent of language. However, we can retain the notion of an unstructured fact. I think this is indeed our most basic notion of “reality”, “the world”, “the way it is”, and this notion can remain invariant through any changes in our concept of “the things that exist”.

IV

In the two passages quoted from Putnam at the outset of this paper he says in the first one that “there are many usable extensions of the notion of an object”, and in the second that “the notions of object and existence have a multitude of different uses rather than one absolute ‘meaning’”. These two charac-

terizations of quantifier variance are subtly different: the first seems to talk about the *possibility* of having different meanings attached to the quantifier, whereas the second seems to say that *in actuality* the quantifier is used with many different meanings. In the preceding discussion I was arguing in behalf of quantifier variance in the first sense, not the second; that is, I was arguing in behalf of *possible quantifier variance*, not *actual quantifier variance*.

In the discussion surrounding the second passage—the passage that seems to describe actual quantifier variance—Putnam criticizes an anti-mereologist who gives the following speech: “I know what you’re talking about if by an object you mean a car, or a bee, or a human being, or a book, or the Eiffel Tower. I even understand it if you refer to my nose or the hood of my car as ‘an object’. But when philosophers say that there is an ‘object’ consisting of *the Eiffel Tower and my nose*, that’s just plain crazy. There simply is no such object. ... and it’s crazy to suppose that every finite universe contains all the objects that those [mereologists] would invent, or, if you please, ‘postulate’. You can’t create objects by ‘postulation’ any more than you can bake a cake by ‘postulation’”.¹⁸

Now I am myself an anti-mereologist who considers the above speech to be quite reasonable. At least it seems reasonable if it is taken in the following spirit: “I assume we’re all speaking plain English, and that we’re employing the quantifier in the sense of plain English. Given that, I understand perfectly well what it means to talk (in plain English) about such things as cars, bees, human beings, books, and the Eiffel Tower, or even to talk about such marginal things as noses and car-hoods. But it’s crazy to say (in plain English) that there exists something composed of my nose and the Eiffel Tower. And you can’t create any such thing by ‘postulating’ it or by changing your language—that would be an absurd form of linguistic idealism”. As a believer in possible quantifier variance I would add: “I can of course make intelligible to myself a possible language that differs from plain English in the meaning of its quantifier. In that imagined language the sentence ‘There exists something that is composed of my nose and the Eiffel Tower’ would be trivially true. If you’re speaking that language for some reason, that’s okay. Just don’t pretend to be speaking plain English”.

Putnam seems to reject this criticism of the mereologist because he thinks that, in plain English, there actually are different senses of the quantifier, and that the mereologist’s claims are true in one of those plain English senses. It seems to me, however, that the linguistic evidence indicates that fluent speakers of English do not speak the mereologist’s language. Speakers of English use what I was calling in previous sections the A-language; they do not use the M-language. My attitude towards the mereologist might be compared to the attitude that I and most current philosophers have towards the traditional epistemologist who says, “Only sense-data can be (strictly) perceived; physical objects cannot be perceived”. Most philosophers would nowadays agree that the sense-data philosopher is saying something that is absurd. This is not to deny that we

can imagine a use of “perceives” different from the one in plain English which would render the sense-data philosopher’s remark true. Indeed, this philosopher’s confusion might be characterized as “language gone on holiday”¹⁹; the philosopher has somehow confused himself into speaking a new language without realizing it. But confusion this is; the philosopher has made a mistake.

What holds for the sense-data philosopher holds as well, I think, for the mereologist. Putnam, it seems, doesn’t see it that way. I’m quite sure he would agree with what I have said about sense-data philosophy, but he apparently views the issue of mereology differently. He says that if an anti-mereologist attacks the mereologist’s view as being obviously wrong, then this shows that the anti-mereologist is a “metaphysical realist” who wrongly rejects Carnap’s “principle of tolerance”.²⁰ I think, however, that the attack can stem from the anti-mereologist’s being a philosopher of common sense.

I think it’s clear that Putnam’s use of “metaphysical realism” has sometimes been confusing. I think he uses this term as the name of a certain attitude in metaphysics as much as the name of a definite philosophical position. To the extent that there is a position it is the denial of even the possibility of quantifier variance. Metaphysical realism says that there is somehow one metaphysically privileged sense of the quantifier, that any departure from this privileged sense would leave us without adequate resources to state the truth properly. There is, so to speak, the quantifier that God would use, and to get things metaphysically right, that’s the quantifier we have to use. I think it is clear, however, that most people whom Putnam would call metaphysical realists have never heard of quantifier variance, and would not readily acknowledge that they have any special views about the necessity to use one kind of quantifier rather than another. It’s their attitude when they philosophize that suggests that they tacitly deny the possibility of quantifier variance. What is the nature of this attitude? Putnam has characterized it in various ways throughout his writings, but I would like to put it as follows: these philosophers argue too much. They descend upon us as a legion of ontological lawyers, their briefcases overflowing with numerous arguments and counterarguments, a case for one entity, a case against another. Questions that appear to be trivial beyond the pale of conversation are somehow converted by them into occasions for deep theoretical debate. “Metaphysical realists” are afflicted with a kind of hyper-theoreticalness. I would certainly not claim that there is an inevitable connection between this affliction and the issue of quantifier variance, but I think that in many cases a potentially helpful diagnosis of the affliction would be roughly as follows: If whenever you make an existential claim in metaphysics you are tacitly or unconsciously assuming that the claim has to be couched in terms of a quantificational apparatus that is in some sense the uniquely right one—the one that God would use—then this assumption is likely to lead you to futile and interminable pseudo-theoretical arguments.

In his discussion of mereology Putnam implies that the proper remedy for the hyper-theoreticalness of “metaphysical realism” is Carnapian tolerance. I

am suggesting that another remedy is “ordinary language philosophy” or an appeal to common sense. Each of these two remedies is appropriate in different cases. Carnapian tolerance is appropriate where an existential sentence being disputed by philosophers is actually vague or ambiguous in plain English (or whatever natural language the philosophers are speaking), and each disputant has in effect become attached to one permissible interpretation of the sentence. In this kind of case once the disputants realize that the quantifier admits of relevantly different interpretations in their language they should each say—thereby exhibiting Carnapian tolerance—that both of them are right, taking the quantifier in the relevantly different meanings. But there are other cases—and I think the case of mereology is an example—in which the disputed sentence admits of only one relevant meaning in plain English, and one of the disputants is saying something that—interpreted in plain English—is trivially absurd. What may prevent this disputant from simply acquiescing to ordinary usage and common sense is the implicit assumption that any question of ontology must be highly theoretical, because any such question turns on how to describe the world in terms of the metaphysically privileged sense of the quantifier. Once the possibility of quantifier variance is accepted, and the notion of a metaphysically privileged quantifier is abandoned, there is nothing to inhibit us from simply expressing the trivial common sense truth in terms of the quantifier we actually have in our language. In these kinds of examples acknowledging quantifier variance leads not to Carnapian tolerance but rather to the common-sense philosopher’s ridicule of needless philosophical paradox.²¹

What may make it difficult for Putnam to see that the mereologist’s position, stated in plain English, seems to be trivially absurd is the fact that there is a closely related position that is not absurd. If we have Clinton’s nose and we have the Eiffel Tower then perhaps we also have the total quantity (or mass) of matter that composes these two things. To talk about *that matter* may seem tolerably correct in plain English. If so, there may be, in terms of plain English, something that can be said to consist of Clinton’s nose and the Eiffel Tower. What needs to be understood, however, is that this is not the thing that mereologist’s mean by “the sum of Clinton’s nose and the Eiffel Tower”. The matter that makes up the nose and the tower at one moment may not be exactly the same matter that makes up these objects at another moment; hence the matter cannot be identified with the mereologist’s “sum” of the nose and the tower. Once this point is grasped I think it seems quite clear that there is no defensible interpretation of the plain English quantifier which would make true the sentence “There exists something that remains composed of Clinton’s nose and the Eiffel Tower while these two objects alter their matter”.

Perhaps the paradoxical effect of the mereologist’s position is clearer if we consider non-contemporaneous objects, such as Socrates’s nose and the Eiffel Tower, which the mereologist will say has a “sum”. The paradox is still clearer when we add to the mereologist’s position another doctrine that almost always goes together with it: the doctrine of temporal parts. If we have temporal parts

and also mereological sums then we wind up with the position, familiar from Quine, Lewis, and many other philosophers, in which any set or series of bits of matter constitutes an object—that is, constitutes an object that can be said to occupy a place and to have ordinary qualities of shape, color, texture, and so on. Imagine that a tree stands in a certain yard, and that nothing extraordinary happens to it. Consider the succession of (stages of) bits of matter that make up the whole tree in the daytime and just the trunk during the nighttime. According to the position just mentioned this succession corresponds to an object on a par with the tree and the trunk. We could then give a name to such an object, say, “a shmree”. A shmree, on this view, is a brown wooden object in the yard that loses its branches every night and regains them every morning. Surely to claim that there is such an object seems to violate common sense to the highest degree. Philosophers who are attracted to this view, I am suggesting, consciously or unconsciously reject the possibility of quantifier variance, and therefore take themselves to be engaged in some high-level theoretical speculation about whether the metaphysically privileged concept of existence—God’s quantifier—would encompass shmrees. We who accept quantifier variance, on the other hand, consider the question about the existence of shmrees to be utterly trivial. It seems immediately clear that our ordinary concept of “a thing” renders the sentence “There is a thing in the yard that keeps gaining and losing branches” false. We can indeed make intelligible to ourselves a different concept of “a thing” which would render the sentence true, but, once we reject as unintelligible the idea of a metaphysically privileged concept of “a thing”, there is no reason for us to resist acknowledging the concept that we in fact have. And, then, there is no reason to resist acknowledging the obvious and trivial truths that flow from that concept.

We can say, in the general spirit of ordinary language philosophy: It’s obviously false to assert that there is a thing in the yard that gains and loses branches, because *that’s not the way we talk*. But the appeal here to “the way we talk” is not to be misconstrued as suggesting the idealist view that, had we talked differently, there would have been something in the yard that gains and loses branches. The suggestion rather is that, had we talked differently, the sentence “There is a thing in the yard that gains and loses branches” would have been true. The appeal to “the way we talk” is a reminder that we have the concepts that we have and that, especially when we do philosophy, we need to remain responsible to the concepts we claim to be using, and not confuse ourselves into slipping inadvertently into a different way of talking.

If the denial of the possibility of quantifier variance leads to a certain kind of hyper-theoreticalness in ontology, the acceptance of quantifier variance leads to an attitude that might be called “ontological shallowness”, or what Quine at one point calls “steadfast laymanship”.²² That is the attitude I am commending. What I have tried to explain is that, depending on what kind of case we are dealing with, ontological shallowness will either take the form of Carnapian

tolerance or the form of insisting that the ontological judgments of common sense be respected in philosophy.

V

I am primarily concerned here with the ontology of physical objects—indeed, highly visible physical objects (roughly, what Austin called “moderate-sized dry goods”²³)—and the issue of quantifier variance with respect to that domain. I’m not addressing questions that involve abstract things (such as the familiar question whether numbers exist in the same sense as physical objects). I have maintained thus far that with respect to the domain of physical objects we can understand possible variations in the quantifier. Are there, however, examples of *actual* variations? Do we, in fact, when describing physical objects in plain English, alter the meaning of the quantifier?

Insofar as English is conceived of as an evolving language, whose major sortal-categories change somewhat from time to time, I’m inclined to think that this change carries with it some change in the meaning of “a thing”, and therefore in the meaning of “the existence (and reference to) a thing”. Moreover, insofar as a given speaker’s repertoire of sortals changes over time, this may qualify as a change in the speaker’s concept of “a thing” (modulo issues about the “communal” nature of each person’s concepts). In any case, the “sortal-dependence” of our concept of “a thing” is an issue I have discussed at length elsewhere, and I will not try to add to it here.²⁴

It’s essential to distinguish between quantifier variance and what David Lewis calls quantifier *restriction*. Lewis points out that if, looking into the refrigerator, I assert, “There is no beer”, I am restricting the quantifier to (the domain of) what is in the refrigerator. The unrestricted quantifier, by contrast, says what exists “strictly” or “*simpliciter*”.²⁵ That the quantifier is often contextually restricted in the way that Lewis says may be granted by all. The question of quantifier variance, however, pertains to the unrestricted quantifier, to our concept of “existence *simpliciter*”. We will see in a moment that Lewis, while often appealing in his work to restrictions on the quantifier, tacitly denies the possibility of there being variations in the unrestricted quantifier.

I have often heard it suggested—though not, to my knowledge, on the basis of anything Lewis himself says—that, once we acknowledge the phenomenon of quantifier restriction, it is no longer clear that the ontological commitments of Lewis and other philosophers, in favor of mereological sums and temporal parts, conflict with the views of common sense. Take, again, the sentence, “There exists a brown wooden thing in the yard that keeps losing its branches every night and gaining them back every morning”. Lewis says that this sentence is true, and ordinary people say it is false. But that’s because Lewis is using the quantifier unrestrictedly, whereas ordinary people are restricting the quantifier to the domain, roughly speaking, of “familiar bodies”. So there is really no conflict, after all.

This suggestion is misguided, I think. Ordinary people must have a concept of “existence *simpliciter*”. They must understand how to use the quantifier unrestrictedly; otherwise no such use could be part of the English language. With respect to the cited sentence, if we explain to ordinary people that the “brown wooden thing” in question need not be any kind of familiar thing, it need not be an interesting thing or the sort of thing one would normally talk about, they still regard the sentence as insanely false, though now qualified in absurdly irrelevant ways. For they take it for granted that there is no brown wooden thing in the yard *of any sort whatever* that keeps gaining and losing branches. They would in fact claim to be able to *see with their eyes* that this is the case. So there surely is the starkest conflict between the views of common sense and the ontological position of Lewis and other philosophers.

I’ve tried to indicate earlier that there is a certain general connection between the tacit rejection of quantifier variance and the tendency amongst many philosophers to reject the ontological judgments of common sense. In Lewis’s case, however, there is a more specific connection between these two rejections. His main reason for believing in mereological sums, contrary (as he admits) to the views of common sense, is that he holds a certain view about quantifier vagueness that, as I will try to explain, comes out of his tacit rejection of quantifier variance.

Let me give an illustration to bring out Lewis’s reasoning about mereological sums. Suppose we are constructing a very simple table out of a top and a leg. If we accept the common sense position, and do not believe in arbitrary mereological sums, we will say that at the beginning of this process, when the top and the leg were far apart (or, perhaps, when no one even had the plan to connect them), there wasn’t any thing made up of the top and the leg. At the end of the process there is a thing, the table, made up of the top and the leg. Now, given the gradualness of this process (the gradualness, say, of driving a nail into a piece of wood or gluing two pieces of wood together), there must surely be a time at which it is indeterminate whether we have a table yet. At that time, then, it is indeterminate whether we have *anything* composed of the top and the leg.²⁶ But that, says Lewis, makes no sense. It can be indeterminate whether there exists something having a certain property only if there exists a thing such that it’s indeterminate whether it has that property. In order for it to be indeterminate whether there exists something composed of the top and the leg, then, there must exist a thing such that it’s indeterminate whether it is composed of the top and the leg. What thing could that be? It seems as if we are being led by the common sense view to say incoherently that there exists a certain thing such that it’s indeterminate whether there exists that thing. “What is this thing such that it sort of is so, and sort of isn’t, that there is any such thing?”²⁷

The only coherent way to view the matter, Lewis concludes, is to assume, contrary to common sense, that at every point in the process there is some-

thing composed of the top and the leg. At some point in the process it may be indeterminate whether the thing composed of the top and the leg qualifies as a table. That's the only kind of indeterminateness that can make sense here. If we restrict the quantifier to tables (or "familiar objects") we can correctly say, "It's indeterminate whether there exists anything composed of the top and the leg", meaning that it's indeterminate whether the sum of the top and the leg is a table (or any other "familiar object".) But the quantifier can be restricted in that manner only against the background of arbitrary mereological sums.

The key assumption in Lewis's argument is the principle mentioned earlier: If it's indeterminate whether there exists something having a certain property, then there exists something such that it's indeterminate whether it has that property. If not for this principle we can simply say that it's indeterminate whether there exists something having the property of being composed of the top and the leg, though there isn't anything such that it's indeterminate whether it's composed of the top and the leg. So why should we accept this principle? If we deny the principle then we are saying that a sentence of the form "There exists something that is *F*" can be vague, not by virtue of the vagueness of the application of the term "*F*" to things, but by virtue of the vagueness of the rest of the sentence, that is, by virtue of the vagueness of the quantifier expression "there exists something". Lewis's argument depends on his denying the possibility that the quantifier can be vague.

Lewis is explicitly aware of this. And he tries to explain why the quantifier can't be vague as follows: "The only intelligible account of vagueness locates it in our thought and language.... Vagueness is semantic indecision. But not all of language is vague. The truth-functional connectives aren't.... Nor are the idioms of quantification, so long as they are unrestricted. How can any of these be vague? What would be the alternatives between which we haven't chosen?"²⁸

I accept Lewis's assumption that vagueness is a matter of semantic indecision.²⁹ I won't say anything here about truth-functional connectives, but what is the problem Lewis is raising about the vagueness of the quantifier? Since the meaning of a logical constant is given by its role in determining the truth-conditions of sentences, the vagueness of the quantifier would consist in our semantic indecision with respect to the truth-conditions of certain sentences, for example, the sentence, "There exists something composed of the top and the leg". With respect to a situation in which the top and the leg are borderline attached, we are undecided about whether to count the sentence as true or false.

I can think of only one reason why Lewis rejects this explanation of the vagueness of the quantifier. The explanation requires us to be able to make sense of there being two possible meanings for the quantifier, the "precisifying" meaning that would make the sentence true and the one that would make it false. Quantifier vagueness requires the possibility of quantifier variance, and Lewis rejects the former because he rejects the latter.³⁰

VI

In this final section I want to try to explain a bit further what I mean by commending a “shallow” approach to ontology. From the standpoint of shallow ontology Quine’s portentous notion of “ontological commitment” already conveys an unfortunate aura of theoretical hype and pseudo-depth. One of the most striking characteristics of “deep ontology” is a certain kind of maddening modesty and caution in its formulations. Whereas the shallow ontologist will address a typical question of ontology either by shrugging it off with Carnapian tolerance for many different answers, or by insisting with Austinian glee that the answer is laughably trivial, the deep ontologist will tend to treat the question with the kind of somber speculative anxiety appropriate to matters of high theory. “All things considered, I am tentatively inclined to be ontologically committed to apple trees but not to apples”.³¹ The challenge for the shallow ontologist posed by this kind of formulation is how to keep a polite straight face while listening to it.

Even from the shallow perspective not every ontological question—within the domain of highly visible physical objects—can be quickly answered. Here is a good philosophical experiment: Look at your hand while you are clenching it, and ask yourself whether some object called a fist has come into existence. As shallow ontologists the first thought that must come to mind when we ask this question is this: *There can’t be anything deep or theoretical here*. The facts are, so to speak, right in front of our eyes. Our task can only be to remind ourselves of relevant ways in which we describe these facts in our language. We might consider, for instance, comparisons and contrasts between how we talk about “making a table” and “making a fist”. (For example, a table can be made and can be destroyed, but do we talk about “destroying a fist”?) Our task is to “*command a clear view of the use of our words*”, as Wittgenstein put it,³² that is, a clear view of how the relevant concepts operate. We’re engaged in what Austin called “linguistic phenomenology”, which need not always be an easy task.³³ Finally—and perhaps most importantly—we are open to the vagueness of “existence”: The best answer might turn out to be that a fist sort of comes into existence and sort of doesn’t.

The deep ontologist approaches this question in a very different spirit. She is engaged in the highly theoretical enterprise of deciding whether her “ontological commitments” should include fists that come into existence when hands are clenched. She will anguish over the “theoretical price” of having such things or not having them. She might eventually venture her best theoretical conjecture about whether fists come into existence. Or perhaps, after furrowing her brow for an appropriate period, she will simply announce that it “darkens the understanding” to suppose that fists keep popping in and out of existence.

Do the deep ontologists care about what common sense says? Sometimes they claim to be trying to balance the demands of common sense against the demands of philosophical theory.³⁴ The trouble is that theory seems always to

win. When I consider the writings of some of the most prominent deep ontologists of recent years—Chisholm, Lewis, van Inwagen, just to mention three—I can rattle off many cases in which they veto a commonsensical judgment in behalf of a philosophical argument, but I would be hard pressed to recall an example in which the reverse happens. My impression is that, in matters of ontology, virtually any theoretical problem, however marginal or flimsy, if it cannot be adequately answered, suffices, by the lights of these philosophers, to trump the most deeply entrenched beliefs of common sense. For a shallow ontologist like me, the opposite is the case. Given any well entrenched ontological judgment of common sense (about highly visible physical objects), I could not imagine giving it up for the sake of some philosophical argument. If I had nothing more definite to say about the argument I would simply repeat Moore's famous point that the force of the common sense judgment shows that there must be something wrong with the argument (even if I don't know what it is).³⁵

I am trying to roughly characterize two ways of approaching an ontological question, what I am calling the "shallow" and "deep" ways. This difference is no doubt partly a matter of intellectual temperament and style, but there is more to it than that. I'm especially interested here in how it relates to the issue of quantifier variance. Let me try to sketch now a general model that shows that common sense judgments of ontology must be taken very seriously. The model, it will be seen, presupposes the possibility of quantifier variance.

Suppose that two philosophers named Xstein and Ystein are engaged in an ontological dispute. I assume that both philosophers claim to be speaking plain English. We might attempt a certain strategy of semantic ascent. Perhaps both philosophers can be gotten to agree that there are two possible languages, call them Xglish and Yglish, such that every sentence asserted by Xstein is true in Xglish and every sentence asserted by Ystein is true in Yglish. Each philosopher agrees that he can understand what the truth-conditions are of the sentences in both languages, and that, in terms of those truth-conditions, he could express in either language everything (every unstructured fact) he wants to assert. In these circumstances, if it were known that Xstein is speaking Xglish and Ystein is speaking Yglish, then these philosophers could simply agree that both are speaking the truth in their respective languages. Of course, this happy result cannot be quite right, since each philosopher claims to be speaking English, which can't be both Xglish and Yglish.

Now suppose, further, that typical fluent speakers of English assert the sentences that Xstein asserts and reject the sentences that Ystein asserts. In other words, they assert sentences that are true in Xglish but false in Yglish. This would seem to be overwhelming evidence that the language they speak is Xglish, not Yglish. As usual, I am focused on sentences having to do with highly visible physical objects. By any reasonable standards of interpretation—for example, the standards explained by Lewis³⁶—we ought to say that these people are uttering truths in Xglish, rather than that they are uttering falsehoods in Yglish.

We have concluded that English is Xglish. We have, therefore, also concluded that Xstein is right and Ystein is wrong. Since they both claim to be speaking English, it is Xstein, not Ystein, who is saying the truth in English. It appears that Ystein has allowed “language to go on holiday”; Ystein’s philosophical theorizing has apparently confused him into in effect slipping into a language of his own creation, the language Yglish. Of the lamentable Ystein we must say, following Wittgenstein, that his theorizing has led him to “have a new conception and interpret it as seeing a new object [or seeing the absence of an old object].”³⁷

As an illustration of this model: Let Lewish be the language in which all of David Lewis’s ontological utterances (with respect to highly visible physical objects) are true; let Inwagish be the language in which all of Peter van Inwagen’s ontological utterances are true; and let Shmenglish be the language in which the typical ontological utterances of non-philosophers are true. (By “true” I always mean “true in the strictest and most philosophical sense”, whatever that is supposed to mean; and I am always assuming an unrestricted use of the quantifier.) The sentence “There exists a brown wooden thing in the yard that keeps gaining and losing branches” is true in Lewish but false in both Inwagish and Shmenglish. The sentence “There exists apples” is true in both Lewish and Shmenglish but false in Inwagish. And so on. I expect the reader to be able to fill in the truth-conditions of the indefinite number of sentences in these three languages.

Both Lewish and Inwagish correspond to Yglish, and Shmenglish corresponds to Xglish. Both Lewis and van Inwagen correspond to Ystein, and I—the shallow ontologist, who simply asserts the same sentences that are asserted by ordinary speakers of English—correspond to Xstein. The argument proceeds as before. Since the typical fluent speaker of English asserts sentences that are true in Shmenglish, but false in both Lewish and Inwagish, the evidence seems to be overwhelming that English is Shmenglish, not either Lewish or Inwagish. Therefore, I am right, and both Lewis and van Inwagen are wrong. Only my ontological sentences are true in English (= Shmenglish); theirs are not. Their sentences are true in either Lewish or Inwagish, but not in English, which is the language they both claim to be speaking.

It will be objected that something essential has been left out of my model. What about the theoretical arguments these philosophers give for their positions? Why haven’t those arguments been taken into account?

But how *can* they be taken into account? That’s my main point here. Those arguments have no obvious bearing on what has now become the critical question: whether English is Shmenglish. Remember, if English is Shmenglish, then it follows directly that both Lewis and van Inwagen are mistaken; that simply follows from the way Shmenglish has been defined. It appears now that all of the deep ontologist’s theoretical maneuvers have no relevance to rebutting the seemingly obvious fact that English is Shmenglish.

I’ve already considered Lewis’s argument about the vagueness of the quantifier. As another and more typical ontological argument, let’s briefly consider

Lewis's reason for believing in temporal parts. Lewis's argument derives from a problem about what he calls "temporary intrinsics".³⁸ The problem, roughly put, is as follows. If an object can have a shape at one moment and lose it at another moment, as common sense says, then the shape is being treated logically and semantically as if it were a relation binding the object to a time. But shapes are intrinsic properties, not relations. The only way to hold onto the intrinsicness of shapes is to claim, contrary to common sense, that what has the shape is a temporal part of the object, and this temporal part never loses the shape.

Now I don't myself see this as a genuinely worrisome problem.³⁹ But I don't want to enter into the details of this issue. My point now is something else. How does Lewis's problem of temporary intrinsics bear on the question whether English is Shmenglish? Remember, once again, that if English is Shmenglish then it follows that Lewis is wrong in asserting (in English) that there are temporal parts. It is given that the following sentence, which Lewis holds to be false in English, but to which typical speakers of English assent, is true in Shmenglish: "A bent wooden stick can sometimes be straightened without any wooden thing going out of existence". In order for the problem of temporary intrinsics to show that English is not Shmenglish the problem would somehow have to be reflected in the speech behavior of typical fluent speakers of English, perhaps by their being disposed to assent to the sentence, "Shapes are intrinsic properties, and an intrinsic property cannot be treated logically and semantically as if it were a relation binding an object to a time". In all likelihood, only typical fluent speakers of English with philosophy degrees from Princeton would have any strong tendency to assent to this sentence. But, furthermore, whatever evidence Lewis might be able to marshal from the speech behavior of typical speakers of English vis-à-vis the problem of temporary intrinsics, this would have to be capable of defeating the overwhelming evidence that English is Shmenglish that derives from the unambiguous disposition of typical speakers to assent to countless sentences such as, "A bent wooden stick can sometimes be straightened without any wooden thing (of any sort whatever, whether familiar or not, interesting or not) going out of existence". I cannot see how Lewis could argue successfully that English is not Shmenglish, let alone argue successfully that English is Lewish.

I am aware that my argument from the Xstein-Ystein model requires various refinements—I intend to provide these elsewhere—but I think that the argument even in its sketchiest form indicates that there is something fundamentally incoherent in the attempt of philosophers to deny basic and seemingly trivial ontological beliefs of common sense.

The point at which my argument depends on quantifier variance is in its assumption that we have these three possible languages, Lewish, Inwagish, and Shmenglish, each with its own quantificational apparatus and concept of "the existence of a thing". A philosopher might try to block the argument by holding that Shmenglish is not a possible language. The burden would fall on

this philosopher to explain why the seemingly most obvious interpretation of the language we speak is not a possibility. I don't see how this burden can be met.⁴⁰

Notes

1. Hilary Putnam, "The Question of Realism" in *Words and Life* (Harvard University Press, Cambridge, Mass., 1994), pp. 304–305.
2. Hilary Putnam, "Truth and Convention: On Davidson's Refutation of Conceptual Relativism", *Dialectica*, Vol. 41 (1987), 69–77, at p. 71.
3. "Truth and Convention", pp. 70 ff.
4. Subject to a slight qualification to be made later.
5. John Wisdom, *Philosophy and Psychoanalysis* (University of California Press, Berkeley, 1969).
6. Ernest Sosa's "existential relativity" has it that things exist relative to conceptual schemes, but Sosa takes great pains to disassociate himself from the idealist claim that language or thought literally creates things. An important question considered in Sosa's discussion, which I'm bypassing here, is how "far down" quantifier variance can go. See Ernest Sosa, "Existential Relativity", *Midwest Studies in Philosophy*, 22 (1999), and my "Sosa's Existential Relativism", forthcoming in the Blackwell volume on *Sosa and His Critics*.
7. So the doctrine need not prevent us from agreeing with Thomas Nagel's claim in *The View from Nowhere* (Oxford University Press, N.Y., 1986), p. 98: "We can speak of 'all the things we can't describe', 'all the things we can't imagine', 'all the things humans can't conceive of', and finally, 'all the things humans are constitutionally incapable of ever conceiving'."
8. I am perhaps making some sort of controversial assumption here related to the inscrutability of reference, or better put, the scrutability of non-reference, for it might be questioned whether some kind of ingenious reference scheme assigning a reference to "touching" might somehow do the trick. I'm quite sure that this is impossible, but perhaps I should limit myself to saying more cautiously that I don't see what such a scheme could be. See, further, my *Dividing Reality* (Oxford University Press, N.Y., 1993), pp. 102–109, and my "Objectivity Without Objects" in the *Proceedings of the Twentieth World Congress of Philosophy: Epistemology*, Vol. 5 (Philosophy Documentation Center, Bowling Green State University, 2000) .
9. There may yet be a sense in which the qualities (features, points of similarity) signified by general words remain the same despite quantifier variance. I am leaving a number of questions open here.
10. This seems to be Putnam's assumption in "Truth and Convention." See Donald Davidson, "On The Very Idea of a Conceptual Scheme", in *Inquiries into Truth and Interpretation* (Oxford University Press, London, 1984).
11. In other words, they have the same "character" in the sense of David Kaplan, "Demonstratives" in J. Almog, J. Perry, and H. K. Wettstein, eds., *Themes from Kaplan* (Oxford University Press, N.Y., 1989).
12. See *Inquiries into Truth and Interpretation*.
13. See Thomas Nagel's criticism of Davidson in *The View from Nowhere*, pp. 94–98.
14. "The Question of Realism", p. 305.

15. See George Bealer, *Quality and Concept* (Oxford University Press, N.Y., 1982), pp. 181–187, and David Lewis, *On the Plurality of Worlds* (Basil Blackwell Ltd, N.Y., 1986), p. 56. It must be borne in mind that I have defined the notion of “truth-conditions”—and consistently use it here—in the coarse-grained unstructured sense.
16. “The Question of Realism”, p. 301.
17. “The Question of Realism”, p. 301 (my emphasis).
18. “Truth and Convention”, p. 72.
19. See Ludwig Wittgenstein, *Philosophical Investigations*, I, 38.
20. “Truth and Convention”, p. 73.
21. I allude casually here to Carnapian tolerance because Putnam mentions it. I am not seriously discussing or defending Carnap’s own thinking on these matters. He seemed to imply that realists ought to “tolerate” phenomenalists, which I think is out of the question for reasons that go beyond the present discussion. See Rudolph Carnap, “Empiricism, Semantics, and Ontology,” in *Meaning and Necessity*, 2nd edition (University of Chicago Press, Chicago, 1956).
22. W. V. Quine, *Word and Object* (MIT Press, Cambridge, Mass., 1960), p. 261.
23. J. L. Austin, *Sense and Sensibilia* (Oxford University Press, London, 1962), p. 8.
24. The notion of sortal-dependence has been introduced to contemporary philosophy by Wiggins’s seminal work, but I cannot tell what his attitude is towards quantifier variance. See David Wiggins, *Sameness and Substance* (Harvard University Press, Cambridge, Mass., 1980). I discuss sortal-dependence in *The Concept of Identity* (Oxford University Press, N.Y., 1982), especially chs. 2 and 3, and “Basic Objects: A Reply to Xu”, *Mind and Language*, Vol. 12, 3 (1997), 406–412.
25. Lewis, *On the Plurality of Worlds*, p. 3.
26. A variation of Lewis’s argument, with convincing arguments that there must be examples of the relevant kind of indeterminateness, is given in Theodore Sider, “Four Dimensionalism”, *The Philosophical Review*, Vol. 106, no. 2, 197–231.
27. Lewis, *On the Plurality of Worlds*, pp. 212–213.
28. *Ibid.*, p. 212.
29. But the argument I am about to go through would be unaffected if we accepted instead Williamson’s epistemic view of vagueness. Lewis and Williamson agree that, for the quantifier to be vague, it must admit of different “interpretations”. The argument is unaffected by whether these interpretations are taken to be different possible “precisifications”, as in Lewis, or different guesses (which we can’t verify) about the one correct interpretation, as in Williamson. See Timothy Williamson, *Vagueness* (Routledge, London, 1994), especially pp. 164, 237, 257–258.
30. An argument similar to Lewis’s, but somewhat more difficult to unravel, stems from the vagueness of identity sentences. Such an argument is given in Sydney Shoemaker, “On What There Are”, *Philosophical Topics*, 16 (1988). I try to show why that argument doesn’t work, appealing in effect to quantifier variance, in my “The Vagueness of Identity”, *Philosophical Topics*, 26 (1999).
31. For the uninitiated, who may possibly think that at least in this example I am tilting at windmills, the example actually expresses the position defended in Peter van Inwagen’s extremely influential *Material Beings* (Cornell University Press, Ithaca, N.Y., 1990).
32. *Philosophical Investigations*, I, 122.

33. J. L. Austin, "A Plea for Excuses", in *Philosophical Papers* (Oxford University Press, London, 1961). It goes without saying that we need not be sympathetic to everything that went under the name "ordinary language philosophy."
34. Lewis often presents his methodology as involving this kind of balancing. See, for example, *Philosophical Papers*, vol. I (Oxford University Press, N.Y., 1983), p. x, and *On the Plurality of Worlds*, pp. 134–135.
35. See G. E. Moore, "A Defense of Common Sense", p. 41, and "Four Forms of Scepticism", p. 222, both in *Philosophical Papers* (Collier Books, N.Y., 1962). See also Ned Markosian, "Brutal Composition", *Philosophical Studies* 92 (1998), 211–249.
36. See his "Radical Interpretation", in *Philosophical Papers*, Vol. I.
37. *Philosophical Investigations*, I, 401.
38. Lewis, *On the Plurality of Worlds*, pp. 202–204.
39. Because to say that a shape is intrinsic might mean intuitively that a thing's having a certain shape at a certain time does not depend on how it is related to other things at that time.
40. In *Dividing Reality* I sought to uncover some general constraints that would make it impossible for a natural language to diverge (too much) from the taken-at-face-value semantic structure of English. My conclusion in that book was that no interesting constraints seem to be defensible. Had the sought-after constraints emerged, although they would certainly have allowed some forms of quantifier variance (in particular, the quantifier variance involved in quantifier vagueness, and that involved in sortal-dependence), they may have implied that either Lewish or Inwagish is not a possible language. An argument for the impossibility of Shmenglish would have to move in the opposite direction: to show that the taken-at-face-value semantic structure of English is not a possibility. In responding to an earlier version of my argument from Shmenglish a few years ago, Sider in fact held that Shmenglish is not a possible language (because, very roughly, its quantifier fails to correspond to the world's logical joints). (See, further, the Introduction to Theodore Sider, *Four Dimensionalism: An Ontology of Persistence and Time* [Oxford University Press, forthcoming].) Will that, I wonder, be the deep ontologist's standard response to my question?