Quine and the Correspondence Theory
Author(s): Hartry Field
Source: *The Philosophical Review*, Vol. 83, No. 2 (Apr., 1974), pp. 200-228
Published by: Duke University Press on behalf of Philosophical Review
Stable URL: http://www.jstor.org/stable/2184138
Accessed: 23/02/2009 09:34

QUINE AND THE CORRESPONDENCE THEORY

IN HIS Dewey lectures (*DL*),[1] Quine states the following doctrine of "ontological relativity":

What makes sense is to say not what the objects of a theory are, absolutely speaking, but how one theory of objects is interpretable or reinterpretable in another [p. 50].

It seems clear from the context that the phrase "objects of a theory" is intended to apply to the objects denoted by the singular terms of the theory and the objects in the extension of the general terms of the theory. So if we say that a predicate "signifies" its extension, we can rewrite the above quotation as:

(1) What makes sense is to say not what the terms of a theory denote or signify, absolutely speaking, but how one theory is interpretable or reinterpretable in another.

This is a very radical contention, for it seems to preclude the possibility of a correspondence theory of truth. By a *correspondence theory* of truth, I mean a theory that says that the notion of truth can be explained by appealing to the relation between words on the one hand and the objects that they are about on the other. The objects that words are about are (by and large) *extralinguistic* objects; so the central feature of a correspondence theory is that it explains truth in terms of some correspondence relations between words and the extralinguistic world. But (1) denies the possibility of such a correspondence theory: it says that the only interesting correspondence you can get is a correspondence between the words of one theory and the words of another.

Quine's only argument for (1) is based on his thesis that semantics is radically indeterminate. But I will argue that *even if semantics is as indeterminate as Quine says it is*, we ought to believe in a correspondence theory of truth and reject (1). If I am right about

[1] "Ontological Relativity," in *Ontological Relativity and Other Essays* (New York, 1969), pp. 26-68.

this and about several of the other points I shall make about indeterminacy, I think it will follow that Quine's radical indeterminacy thesis is of considerably less philosophical interest than is usually supposed.

Quine's main argument for the claim that semantics is radically indeterminate is physicalistic. At the beginning of *DL* he writes:

> [K]nowledge, mind, and meaning are part of the same world that they have to do with, and . . . are to be studied in the same empirical spirit that animates natural science. There is no place for a priori philosophy [p. 26].

He then goes on to suggest that once this position is taken seriously, one is bound to recognize the existence of indeterminacy. Suppose, for instance, that we are interested in determining the extension of the foreign term "*gavagai*." If we look at the matter physicalistically, we see that there is no sense in saying that "*gavagai*" has the set of rabbits as its extension as opposed to the set of undetached rabbit parts, unless we can find physical facts—facts about the speaker's behavioral dispositions, his causal relations to rabbits, and so on[2]—which determine that it is the set of rabbits rather than the set of undetached rabbit parts that is the real extension of the term. And Quine thinks it is obvious that there *are* no physical facts underlying the use of the term that could allow us to say that the term signifies one of these sets rather than the other.

To set out the matter in a bit more detail, let us suppose that "*gavagai*" and "*potrzebie*" are foreign terms that are most naturally translated as "rabbit" and "dinosaur," and that "*glub*" is a term of the same language that is most naturally translated as "is identical to." Then according to Quine, there is no fact of the matter as to whether

(*i*) "*gavagai*" signifies the set of rabbits, "*potrzebie*" the set of dinosaurs, and "*glub*" the identity relation; or

(*ii*) "*gavagai*" signifies the set of undetached rabbit parts,

[2] Quine thinks that it is sufficient to consider behavioral dispositions alone; but to a large extent his discussion can be freed from this dubious assumption (and from his equally dubious verificationist assumptions), as the next paragraph illustrates.

> "*potrzebie*" the set of undetached dinosaur parts, and "*glub*" the paridentity relation (that is, the relation of being undetached parts of the same object).

To see why this seems plausible to Quine, suppose that there *were* a fact of the matter; suppose, for instance, that (*i*) were really true and (*ii*) really false. Then it ought to be possible to state facts about the way that "*gavagai*" is used which *make* this word a word for rabbits rather than for undetached rabbit parts. One place we might look for such facts is in the causal links between the rabbits on the one hand and the uses of "*gavagai*" on the other. But this does not seem to work: any causal links between *rabbits* and uses of "*gavagai*" are also causal links between *undetached rabbit parts* and uses of "*gavagai*," so it appears that causal connections by themselves will not do the job. What then are we to supplement them with? Perhaps the foreigners' dispositions to assent and dissent? Suppose that when we place a foreigner in an environment containing exactly one visible rabbit, he tells us that for every pair of nearby *gavagais* x and y, x *glub* y. This fact about the foreign speakers determines that *if* "*glub*" is a word for identity *then* "*gavagai*" is a word for rabbits, and that *if* "*glub*" is a word for paridentity *then* "*gavagai*" is a word for undetached rabbit parts. But what facts determine whether "*glub*" is a word for identity or paridentity? There is no obvious answer to this question that does not assume an answer to the question of whether "*gavagai*" or some similar word is a word for "whole objects" or for their undetached parts. Again the attempt to find physical facts which decide between (*i*) and (*ii*) fails, and it is hard to see where else such physical facts are to be found. Someone might perhaps claim that either (*i*) is right and (*ii*) wrong, or vice versa, *even though no physical facts determine which*; but this is the position that Quine calls "the myth of the museum" and which he rejects on physicalistic grounds.

I think that Quine is correct in holding that "the myth of the museum" is a totally unreasonable position: if indeterminacy is to be rejected, it must be rejected by finding physical facts which do in some sense decide between (*i*) and (*ii*). I believe that there *are* physical facts which (in the relevant sense) decide between (*i*) and (*ii*), but I shall say nothing about that in this paper. Instead, I shall pretend to believe that Quine is right about this example of indeterminacy, and consider the consequences that such indeterminacy would have for the correspondence theory of truth.

The first thing I want to do is note that if such indeterminacies exist, then at least *half* of (1) is correct: it makes no sense to ask what the terms of a language or theory refer to (denote or signify). Actually, what I have just said goes beyond (1) in some respects, for what Quine claims in (1) is that it does not make sense to ask what the terms of a language or theory refer to "absolutely speaking," while my remark in the previous sentence suggests that (if the indeterminacy thesis is right) it does not make sense to ask what terms refer to *either* in the absolute sense *or* in Quine's relativized sense. For the moment, however, let us ignore the notion of relativized reference, and use the terms "refer,""denote," and "signify" in the ordinary, absolute way. In that case, since there is no fact of the matter as to *what* a term like "*gavagai*" signifies (has for its extension), it seems pretty obvious that there is no sense in speaking of "the extension" of the term.

This conclusion may seem to rule out the possibility of a correspondence theory of truth, for the most obvious form of the correspondence theory of truth is one which explains truth in terms of such "correspondence relations" as denotation and signification.[3] What I claim, however, is that a correspondence theory is still possible: all that is needed is the introduction of certain *more general* correspondence relations between words and extralinguistic objects (or sets of objects). For instance, the difficulty with the relation of signification was that we had to choose *between* saying that "*gavagai*" signified the set of rabbits and saying that "*gavagai*" signified the set of undetached rabbit parts, and that

[3] Such a correspondence theory would say, for instance, that for "Caesar crossed the Rubicon" to be true there must be objects x and y and a relation (in extension) R such that "Caesar" denotes x, "the Rubicon" denotes y, "crossed" signifies R, and x bears R to y. This illustrates how such a correspondence theory would work for relatively simple sentences; and by employing Tarski's work on truth we can easily extend the treatment to more complex sentences. (See my paper, "Tarski's Theory of Truth," *Journal of Philosophy*, LXIX [July 13, 1972].)

according to the indeterminacy thesis there is no physical basis for such a choice. So to avoid having to make such a choice, why not introduce a new correspondence relation—call it "partial signification"—and say that the term "*gavagai*" bears this correspondence relation *both* to the set of rabbits *and* to the set of undetached rabbit parts? (Each of these sets will then be called *partial extensions* of the term "*gavagai*": so even though the term has no extension, it has a number of different partial extensions.) Of course, the introduction of new correspondence relations like partial signification is not too interesting unless it proves possible to explain truth in terms of them; this raises problems which I will turn to shortly. But let us ignore these problems of detail for the moment, and say more about the general idea of handling indeterminacy by introducing new correspondence relations.

To see the import of this idea, and how the idea differs from Quine's ideas, let us shift our attention from foreign languages to our own language—or, better yet, to my own language as used by me right now. It is clear that the indeterminacy argument outlined before applies just as cogently to our own language as it does to the foreign language. In other words, we can argue that there is no fact of the matter as to whether

(*i**) "rabbit" (as I use it now) signifies the set of rabbits, "dinosaur" the set of dinosaurs, and "is identical to" the identity relation; or

(*ii**) "rabbit" (as I use it now) signifies the set of undetached rabbit parts, "dinosaur" the set of undetached dinosaur parts, and "is identical to" the paridentity relation,

on the grounds that there are no physical facts that could decide between (*i**) and (*ii**). To say, in the face of this argument, that (*i**) is really true and (*ii**) is really false is to fall victim to the myth of the museum.[4] So I think that if we believe in this example

[4] Note that this argument shows that the "disquotation schema"
"——" signifies the set of ——'s and nothing else
cannot be accepted by anyone who believes in indeterminacy. The schema does provide a partial axiomatization of the concept of signification; but to axiomatize a concept is not to show that the concept is physicalistically acceptable. (For more on physicalistic acceptability, see Secs. III–V of my paper on Tarski.)

of indeterminacy, we must give up speaking of the extension of "rabbit," and say instead that "rabbit" (as we use it right now) has the set of rabbits as one partial extension and the set of undetached rabbit parts as another.[5]

The point of view I have just outlined is very different from Quine's. On Quine's view, there is no need to *give up* the ordinary semantic notions of denotation and signification (or extension); instead, we can *relativize* them. Consider first the case where we are giving a semantics for a foreign language. On Quine's view there is no need to abandon all talk of what a foreign term like "*gavagai*" signifies: what we must do, however, is say that *relative to the obvious translation manual* it signifies the set of rabbits, and

---

It does not follow that the believer in indeterminacy is deprived of *all* disquotation schemas: he may be able to adhere to *unusual* schemas like

"——" partially signifies the set of ——'s and partially signifies the set of undetached —— parts.

Such schemas might axiomatize, and clarify, the concept of partial signification to precisely the same degree that the more usual schemas axiomatize, and clarify, the more usual semantic concepts.

[5] The concept of partial signification is not really as unfamiliar as it may sound, for we implicitly employ it in giving the semantics of *vague* expressions. Suppose we were asked what the extension of the English phrase "tall man" is—is it the set of men taller than 6′ 0″, or the set of men taller than 6′ ½″, or what? Clearly there is no fact of the matter as to which of these sets is "the real extension" of the English phrase "tall man"; for clearly "tall man" does not simply signify a particular set, but *partially* signifies various different sets—viz., those sets of form

$\{x \mid x \text{ is a person whose height is greater than } h\}$

where h takes on values in some region centering around six feet and extending a few inches in either direction. (A still better account of vagueness could be given by quantifying the notion of partial signification: thus—if we pretend that the set of possible heights is discrete—we could say that "tall man" signifies various different sets to various different degrees.)

That the notions of partial denotation and partial signification (or quantified forms of these notions) are of use for dealing with vague expressions is by no means a novel point; e.g., it is suggested by David Lewis in the appendix to "General Semantics" (*Synthèse*, 22 [1970], 18-67). These cases are unexciting because they are very unsurprising, and because we can easily do without such vague terms whenever our conversation turns to exact and serious purposes. But there is at least an abstract possibility that something similar to vagueness occurs even where we least expect it: it is possible that there are much more deep and pervasive ways in which our terms have indeterminate application to the world. And what Quine's argument about "*gavagai*" seems to me to suggest is that this abstract possibility is in fact realized.

that *relative to an unobvious but nonetheless acceptable manual* it signifies the set of undetached rabbit parts. The central role that translation manuals play in Quine's semantics reflects the doctrine of ontological relativity: the view that it makes no sense to speak of "absolute" correspondence relations between words and extra-linguistic objects, and that what does make sense is to say how one language or theory is translatable and retranslatable into another.

There is, however, a serious difficulty with Quine's view: the notion of a general term signifying a set *relative to a given translation manual* (or of a singular term denoting something *relative to a given manual*) does not seem to make any sense. (1) seems to suggest that we can explain the idea of signification relative to a manual as follows:

> (2) To say that a term T used in one language signifies the set of rabbits, relative to a translation manual M, is in effect just to say that M translates T as "rabbit."

But this is not a satisfactory explanation. The difficulty becomes clear when we try to define explicitly the notion of relative signification on the model that (2) suggests. The obvious first attempt is

> (3) For every predicate T, set $\{x \mid Fx\}$, and manual M, T signifies $\{x \mid Fx\}$ relative to M if and only if M maps T into "F."

But this clearly involves a use-mention confusion since we are trying to quantify over a variable "F" that appears both inside and outside quotation marks.[6] Can (3) be modified so as to avoid this defect? Yes: we can say

> (4) For every predicate T, set y [or $\{x \mid Fx\}$], and manual M, T signifies y [or $\{x \mid Fx\}$] relative to M if and only if M maps T into some term which signifies y [or $\{x \mid Fx\}$].

---

[6] The point would perhaps be a bit clearer for singular terms than for general terms like "*gavagai*," since certain irrelevant issues about the use of predicate letters as variables would not then arise. For singular terms, the unacceptable "definition" analogous to (3) is

(3′) For every name T, object x, and manual M, T denotes x relative to M if and only if M translates T as "x."

But (4) defines Quine's relativized notion of signification only in terms of an unrelativized notion of signification applied to our own language (the language into which we translate). So it seems that *we have to understand this unrelativized notion before we can understand the relativized notion employed in (2).*

The difficulty is obvious: the whole point of relativizing the notions of denotation and signification to a translation manual was that due to the indeterminacy of reference (or "inscrutability of reference," to use Quine's phrase) the unrelativized notions of denotation and signification are not physicalistically acceptable. But the foregoing remarks show that once this indeterminacy is taken seriously and applied to our own current language as well as to other languages, the manual-relative notions of denotation and signification are not acceptable, either. By employing them, Quine himself has become a victim of "the myth of the museum."

It is clear from these remarks that what Quine needs for his notion of reference-relative-to-a-translation-manual to make sense is some link between the word "rabbit" of our own language and the actual rabbits. But the problem is that the indeterminacy thesis denies the existence of any such connection: it denies the existence of any connection between "rabbit" and *rabbits* that does not also hold between "rabbit" and *undetached rabbit parts.* Quine tries to evade this problem in *DL* by introducing a new sort of relativity—relativity to a background language. He admits that we cannot say that "rabbit" (as we use it) refers *absolutely* to the rabbits; but, he suggests, surely no one can complain if we say that "rabbit" refers to rabbits relative to our own language? Thus he writes:

> It is meaningless to ask whether, in general, our terms "rabbit," "rabbit part," "number," etc., really refer respectively to rabbits, rabbit parts, numbers, etc., rather than to some ingeniously permuted denotations. It is meaningless to ask this absolutely; we can meaningfully ask it only relative to some background language. . . . Querying reference in any more absolute way would be like asking absolute position, or absolute velocity, rather than position or velocity relative to a given frame of reference [*DL*, pp. 48-49].

Unfortunately, Quine's suggestion will not work: if the indeterminacy thesis precludes us from making sense of any relation between "rabbit" and the rabbits that is not equally much a connection between "rabbit" and the undetached rabbit parts, it is clear that merely by relativizing to our language (that is, to our word "rabbit" and to a whole bunch of other expressions) there is no hope of improving the situation.[7] It is clear, then, that Quine's indeterminacy thesis forces us to give up not only the

---

[7] I have not denied that we could make sense of the question of what English terms (if any) a foreign term is *codenotational with* or *coextensive with* (relative to a manual M): I have denied only that we can make sense of the question of what extralinguistic objects or sets of objects a foreign term denotes or signifies (relative to a manual *M and to English*).

In other words, my argument does not directly refute doctrine (1)—though it does show that the phrase "absolutely speaking" which occurs in (1) is misleading and should be dropped. What it does do is to raise the cost of doctrine (1) higher than Quine or anyone else is likely to be willing to pay. For the notions of codenotationality and coextensivity are not powerful enough semantic notions to be of much utility—e.g., they are of no use whatever in a theory of truth, for it is impossible to define truth (or even material equivalence) in terms of them.

Quine does not accept the radical proposal that we make do with the notions of codenotationality and coextensiveness: he thinks he can reject this proposal and still adhere to (1). This seems plausible to him because of his well-known analogy between the doctrine of ontological relativity and Leibniz' relational doctrine of space. Leibniz held—in analogy to (1)—that it makes no sense to speak of relations between physical objects and absolute space, and that what does make sense is to speak of spatial relations between physical objects. But Leibniz held that locutions which appear to state relations between physical objects and absolute space (e.g., "object x has position p or velocity v") do not need to be *given up*; we can reinterpret them so that they do not state such relations, by relativizing them to a co-ordinate system composed of physical objects. Similarly, Quine suggests, we can adhere to (1) without giving up the predicates of denotation and signification, as long as we relativize them to a "co-ordinate system" of words.

But there is a crucial disanalogy here: on Leibniz' theory, we can understand relativized claims about the relations of physical objects to places *only because places are understood as constituted by the relations of physical objects*; whereas no one holds that physical objects are constituted by the relations of words. This difference leads to a further difference of more direct relevance: whereas the relativized predicate "x has velocity v relative to y" is definable in terms of the spatial relations between x and y (viz., as the time derivative of the distance), the relativized predicate "T_1 denotes x relative to T_2" is not definable in terms of the linguistic relations between T_1 and T_2. In fact (and this is my argument in the text), there is no hope of defining this last predicate at all unless we can establish a relation between either T_1 or T_2 on the one hand, and x (and no object other than x) on the other. But that is just what the indeterminacy thesis precludes us from doing.

absolute notions of denotation and signification, but even the relativized notions which Quine has proposed as surrogates for them. (See note 7 for a further discussion of this.)

I have urged, however, that there is nothing in the indeterminacy thesis which forces us to give up various *generalizations* of the notions of denotation and signification—for example, the notions of partial denotation and partial signification, which are just like the notions of denotation and signification except that a single term is allowed to partially denote more than one object or to partially signify more than one set. Still, the existence of such correspondence relations as these is not a cause for much satisfaction unless we can use them in an explanation of truth and falsity. What I want to do is investigate the question of how this can be done.

In doing this, it will be convenient to start with a simpler example of indeterminacy. (The example is one which I have employed elsewhere; and I believe that it is a *genuine* example of indeterminacy, though it is an indeterminacy of a much more limited variety than the indeterminacy which Quine's examples purport to establish.) I claim that we can translate certain outdated physical theories into current theory in a variety of ways; no one translation is best. Suppose, for instance, that we want to translate Newtonian mechanics into special relativity. Then there are two natural ways to translate the word "mass": we could translate it as "relativistic mass," or we could translate it as "rest mass." If we translate it in the first way, Newton's tenet

(5) Momentum is mass times velocity

comes out strictly true, but his tenet

(6) Mass is invariant (that is, independent of the frame of reference)

comes out false (though approximately true at low velocities). If, on the other hand, we translate it in the second way, (6) comes out strictly true and (5) comes out false (but approximately true at low velocities). I claim that there is no fact of the matter as to which of these translations is "the correct one," and hence no fact of the matter as to which of Newton's sentences (5) and (6) was strictly true.[8] I will not argue for this contention here, since I have done so elsewhere[9] and since the details of the argument are not too important for my present purposes.

I hope that my earlier remarks have made clear the sort of way I want to treat such examples of indeterminacy: on my view, what the example shows is that Newton's word "mass" partially denoted both relativistic mass and rest mass; since it *partially* denoted *both* these quantities, it did not *fully* (or *determinately*) denote either. Perhaps I should explicitly mention that the notion of partial denotation is meant to apply to terms that have determinate application as well as to terms that do not: a term with determinate application is a term that partially denotes exactly one thing. So when we give a semantics for sentences that include both determinate and indeterminate expressions, there is no need to employ the concept of denotation *in addition to* the notion of partial denotation: the latter concept is just a generalization of the former.

Let us now provisionally sketch such a semantics. The first step is to introduce the model-theoretic notion of a structure. For our purposes, we can say that a *structure* for a language L is a function that maps all the names of L into extralinguistic objects and all the predicates of L into sets of extralinguistic objects. Note that a language will have many structures: for instance, there are structures which assign Ted Williams to Newton's term "mass," and there are other structures which assign the set of aardvarks to "is invariant." But if we ignore the existence of indeterminacy for a moment, there is one structure for the language L that is worth singling out—namely, the one that assigns to each name of L the object that name denotes, and that assigns to each predicate of L the set of objects that that predicate signifies. Let

[8] The "mass" example illustrates Quine's contention that two acceptable manuals for translating a foreign language or theory might "dictate, in countless cases, utterly disparate translations; not merely mutual paraphrases, but translations each of which would be excluded by the other system of translation. *Two such translations might even be patently contrary in truth value*, provided there is no stimulation that would encourage assent to either" (*Word and Object* [Cambridge, 1960], pp. 73-74; italics mine).

[9] "Theory Change and the Indeterminacy of Reference," *Journal of Philosophy*, LXX (August, 1973).

us say that this structure is the one that *accords with* the semantics of L.

Once we remember the existence of indeterminacy, however, we lose the ability to single out a unique structure in this way. But we can do the next best thing: we can introduce a *class* of structures, each of which *partially accords* with the semantics of L. We can provisionally define this notion of "partially according" as follows:

> (7) A structure m partially accords with the semantics of L if and only if each term of L partially denotes or partially signifies the entity which m assigns to it.

Partial accordance is of course just a generalization of accordance; when all the terms of L are determinate, only one structure partially accords with the semantics of L—namely, the one that fully accords with the semantics of L.

Now that we have the notion of a structure, it is easy to explain the notion of truth-in-a-structure: to say that a sentence is true in the structure m is in effect to say that it *would* be true if all the terms in the sentence were determinate and if they denoted or signified just those entities which m assigns to them. Of course, this definition of truth-in-a-structure is only a vague and intuitive one; but anyone familiar with Tarski's work on truth will know that it can be made perfectly clear and precise. For present purposes, it will be enough to illustrate the definition by the following example: relative to a structure that assigns Ted Williams to the English word "mass" and the set of aardvarks to "is invariant," the English sentence "Mass is invariant" comes out untrue (since Ted Williams is not an aardvark).

This example makes it clear that the notion of truth-in-a-structure is not of much interest for *every* structure: the fact that Newton's sentence (6) comes out untrue-in-the-structure-just-considered has no bearing whatever on whether it was true in any ordinary sense of "true." It is obvious, then, that we must somehow restrict our attention to those manuals which are in close accord with the semantics of the language. When all the terms of L are determinate, so that there is a unique structure which perfectly accords with L, it is clear enough how this should be done: we should say that a sentence is true if and only if it is true in the unique structure that accords with *L*. But what do we do when there are indeterminate terms and hence no uniquely privileged structure? Obviously, we want somehow to restrict our attention to those manuals which partially accord with the semantics of *L*. But there are two very different ways in which this "restriction of attention" can be accomplished.

The most natural way to do it is to define "true" in terms of "true-in-*m*" and "partially accords," as follows:

> (8) A sentence of *L* is true if and only if it is true-relative-to-*m* for every structure *m* that partially accords with the semantics of *L*.

This would have the consequence that when Newton uttered the disjunction of (5) and (6), what he said was true, and that when he uttered their conjunction what he said was false (where to say that a sentence is false is to say that its negation is true). It would also have the consequence that when Newton uttered (5) by itself, or (6) by itself, what he was saying was untrue (and also unfalse, if "false" is defined as above). It seems to me that this is precisely the conclusion we want. For it would be unreasonable to say that (5) was true and (6) was untrue: this would suggest that there was a fact of the matter, that Newton's term was really a word for relativistic mass. It would be equally unreasonable to say that (6) was true and that (5) was untrue: again that would suggest that there was a fact of the matter as to which conjunct was true. So unless we are to say that *both* (5) and (6) were true, which is clearly unreasonable,[10] the only alternative is to say that *neither* was true. (8) lets us say just this, without forcing us to give up the view that the disjunction of (5) and (6) was true and the conjunction false.

In spite of what I have said in defense of (8), philosophers of a more Quinean persuasion may prefer to handle the matter differently. The alternative they may propose is to make an arbitrary

[10] A set of true sentences ought to have only true consequences, and yet (5) and (6) together have the false consequence that momentum divided by velocity is invariant.

choice among all of the structures which partially accord with the semantics of L, and to let "true" stand for "true relative to m_o" where "m_o" denotes the structure arbitrarily chosen. This procedure would have the consequence that the pair consisting of (5) and (6) contains one true sentence and one false sentence; though which one is true and which one false would depend on which structure is arbitrarily chosen. This proposal seems to me to be completely unreasonable. A semantics according to which (5) was true and (6) false seems to me to be a semantics committed to the idea that "mass" denoted relativistic mass, and a semantics according to which (6) was true and (5) false seems to me to be a semantics committed to the idea that "mass" denoted rest mass. To make an arbitrary choice between these theories is to make an arbitrary choice between two theories *each of which is inadequate* (given the existence of indeterminacy). For this reason I think we should reject such arbitrary choices and adhere to the definition (8).

Perhaps it is worth mentioning another respect in which the "arbitrary choice" approach is inferior to (8). The difficulty arises from the fact that determining which structures partially accord with a particular language L and which ones do not is a matter of empirical linguistics, and it is easy to make errors in the process (especially in the case of foreign languages). This creates a difficulty for the "arbitrary choice" approach, because the only way for a person to make a choice is for him to choose among the structures he *thinks* to be in partial accordance with the language, so that in choosing he may end up with a structure m_o that does not accord with the language at all. Suppose this happens. If in this situation he adheres to the "arbitrary choice" approach by letting "true" be an abbreviation of "true relative to m_o," then a sentence might be "true" in his sense even though it is false-relative-to-m for every manual m that partially accords with the language. This result seems extremely unpalatable, and I know of no way to avoid it other than to forgo arbitrary choices and to adhere to definition (8).

Let us now look at what we have accomplished. If we combine (7) and (8) with Tarski's definition of truth-in-a-structure, the result is a definition of truth in terms of partial denotation and partial signification which I will call the *restricted truth-definition*. The restricted truth-definition is a generalization of the usual variety of truth-definition, which is *highly* restricted: the highly restricted truth-definition defines truth in terms of denotation and signification, and hence works only for languages containing no indeterminate terms (if there are any such languages); while the restricted truth-definition can handle not only those languages but also languages with expressions that are indeterminate in the way that Newton's term "mass" was indeterminate. But even the restricted truth-definition cannot adequately handle a phenomenon that I call *correlative indeterminacy*. Correlative indeterminacy is an all-pervasive phenomenon, if Quine is right, for the "*gavagai*" example is an example of correlative indeterminacy if it is an example of indeterminacy at all.

To see what I mean by correlative indeterminacy, let us consider a foreign sentence that contains both "*gavagai*" and "*glub*," together with various logical connectives which we will assume to have a determinate semantics. Actually, there is no need to focus on foreign words, for since we have already rejected the translation-theoretic approach to semantics, we can just as well give the semantics directly for English. Consider then the sentence

(9) $\forall x \forall y$(x and y are nearby rabbits $\supset$ x is identical to y),

which we imagine to be uttered in the environment of exactly one rabbit. Obviously, (9) ought to come out true in this environment, but if we try to apply the restricted truth-definition to it we do not get this desired result. The reason is that, by our earlier suppositions, "rabbit" partially signifies both the set of rabbits and the set of undetached rabbit parts, and "is identical to" partially signifies both identity and paridentity. But then (7) has the consequence that not only are the following two structures relevant:

(*a*) "rabbit" $\rightarrow$ {rabbits},
"is identical to" $\rightarrow$ identity

(*b*) "rabbit" $\rightarrow$ {undetached rabbit parts},
"is identical to" $\rightarrow$ paridentity

(where a *relevant* structure is one which partially accords with the semantics of L); it also has the consequence that the following two "undesirable" structures are relevant:

(*c*) "rabbit" $\rightarrow$ {rabbits},
"is identical to" $\rightarrow$ paridentity

(*d*) "rabbit" $\rightarrow$ {undetached rabbit parts},
"is identical to" $\rightarrow$ identity.

It is (*d*) that causes particular problems for (9): since not all the undetached rabbit parts are identical to each other, (9) comes out false in some relevant structures, and therefore by (8) it cannot be true.

It is clear that no advantage is to be gained by abandoning (8) and going back to the "arbitrary choice" approach to indeterminacy: obviously what needs revision is (7). But it is equally obvious that there is no possibility of making the required modification of (7), unless we introduce new correspondence relations besides partial denotation and partial signification. (I emphasize that these must be correspondence relations as that term was explained in my introductory remarks: that is, they must be relations which hold between words and extralinguistic entities rather than between the words of one language and the words of another.) What we need to do, then, is introduce slightly more complicated correspondence relations than the ones we have so far, and then define relevance in terms of them. The new definition, when combined with (8) and with Tarski's definition of truth-in-a-structure, will give us an *unrestricted truth-definition* that explains truth in terms of our new correspondence relations.

In order to see what new correspondence relations we need, let us return to a consideration of the details of the "rabbit" example, as they were set out at the beginning of this paper. Let us focus first not on the predicate "rabbit," but on the predicate "is identical with." A philosopher who did not believe in the existence of indeterminacy would presumably want to say that this predicate signified identity—that is, that it signified the set of those ordered pairs whose first member and last member were the same. The problem that Quine has raised for such a philosopher, however, is to say *in virtue of what* this predicate signifies that relation. The most obvious facts to cite are that people use the predicate "is identical with" in such a way that it obeys certain laws: the laws governing equivalence relations, and the substitutivity schema, "$\forall x \forall y (x = y \land Fx \supset Fy)$" (restricted to those contexts we call extensional). But the fact that people use the predicate "is identical with" in accord with these laws is not sufficient to rule out the hypothesis that it signifies paridentity—at least, not without the assumption that the language contains predicates true of one object and false of a paridentical object, and it seems that we can always refuse to grant this assumption (for example, by saying that "white" is true of the undetached parts of white things, rather than being true of the white things themselves). So there is certainly a prima-facie difficulty here in assuming that "is identical with" is a word for identity rather than paridentity, and unless this difficulty can be resolved there is little to say except that it partially signifies both.

Now let us turn to the words "rabbit" and "dinosaur." It seems fairly clear that there are facts about the causal relations (for example, perceptual relations) between rabbits and their undetached parts on the one hand and our uses of "rabbit" on the other by which we could hope to explain that "rabbit" was a word for some kind of rabbitish entity (whether the rabbit itself, or its undetached part, or whatever); and it seems equally clear that there are facts about the causal relations between dinosaurs and ourselves (for example, our perceptual relations to dinosaur *fossils*) with which we could hope to explain that "dinosaur" was a word for *some* dinosaurish kind of entity. The only problem is to figure out precisely *which* kind of entity. We get a certain amount of help here from the circumstances under which people are disposed to assent to (9): this tells us that *if* "is identical with" were a word for identity, then "rabbit" would signify the set of rabbits, and that *if* "is identical with" were a word for paridentity, then "rabbit" would signify the set of undetached rabbit parts. But the indeterminacy thesis forces us to assume that "is identical with" is not a word for either identity or paridentity alone: it is a word that partially signifies both. What then are we to say about "rabbit"? One thing we could say is that it partially signifies the set of rabbits and partially signifies the set of undetached rabbit parts. But we can also say something a bit more informative: we can say that relative to a correlation of identity with "is identical to," "rabbit" signifies the set of rabbits; and that relative to a correlation of paridentity with "is identical to," "rabbit" signifies the set of undetached rabbit parts. What I am suggesting, then, is that if we take Quine's radical indeterminacy thesis seriously we should take "rabbit," "dinosaur," and so forth as being *dependent* predicates, predicates whose extension is a function of the extension of another predicate, "identical" (which I will call the *basis* of the dependent predicates). Then if this basis predicate turns out to be indeterminate, the other predicates which functionally depend on them will turn out indeterminate too (that is, they too will partially signify more than one set); but the functional dependence of these other predicates on the basis predicate "identical" will allow us to correlate the partial extensions of one predicate with the partial extensions of others. This means that in working out the key step of our unrestricted truth-definition—namely, the analogue of (7)—we will be able to toss out those structures which assign uncorrelated partial extensions to our predicates.

The above paragraph should make clear how the new definition of "partially according" is going to work. A central concept we must employ is that of one term t_1 being the basis of another term t_2 (which I will write as "$t_1 = b(t_2)$"). The notion of a name or predicate being dependent is of course definable from this: a name or predicate is dependent if it has a basis. In the Quinean example lately considered, the dependent terms "rabbit" and "dinosaur" had as their basis the independent term "identical." We do not have to require in general that the basis always be independent, but we do have to require that if a term has a dependent term as its basis, then either the basis of the basis is independent, or the basis of *that* is independent, or . . . . Call this the *grounding requirement.*

We can now define relevance as follows:

(7*) A structure m partially accords with the semantics of L if and only if

(*a*) each independent term t of L partially denotes or partially signifies $m(t)$;

(*b*) each dependent term t of L denotes or signifies $m(t)$ relative to the correlation of $m(b(t))$ with $b(t)$.

If the grounding requirement is met, (7*) is a formally acceptable recursive definition, and combined with (8) and with Tarski's definition of truth-in-a-structure it gives us the definition of truth that we wanted.[11]

The treatment of indeterminacy that I have been advocating raises a number of epistemological questions: for example, how do we know which things a term partially refers to; and how do we know which terms have bases and what their bases are? I cannot adequately discuss these epistemological questions in this paper, but perhaps a few observations are in order.

One important factor to take into consideration in deciding what partially denotes what, what is the basis of what, and so on, is the consequences that such decisions have for the truth and falsity of sentences. The main reasons for inventing the notions of denotation, partial denotation, basis, and so forth, was to give us a reasonable theory of truth, and if our decisions as to what partially denotes what (and so forth) lead us to unreasonable conclusions about truth then those decisions are inadequate.

But this requirement (which I label "Requirement 1") is not a sufficient one: we also want the semantic notions we employ to

[11] Added generality could be obtained by replacing "denotes" and "signifies" in clause (*b*) of (7*) by "partially denotes" and "partially signifies"; and it is easily seen that this added generality is necessary if we are to handle "mass"-type indeterminacy and "*gavagai*"-type indeterminacy together. (Consider especially "mass"-type indeterminacy for *predicates*.) A second way to generalize (7*) is to allow dependent terms to have more than one basis; a third way is suggested in the ensuing paragraphs on Benacerraf's problem. Putting the three generalizations together, we get

(7**) A structure m partially accords with the semantics of L if and only if

(*a*) each independent atom t of L partially denotes or partially signifies $m(t)$.

(*b*) each dependent atom t of L partially denotes or partially signifies $m(t)$, relative to the correlation of $m(b_1(t)),\ldots, m(b_n(t))$ with $b_1(t),\ldots, b_n(t)$, respectively.

I think that this definition is general enough to handle every plausible example of indeterminacy. There are some apparent exceptions to this claim—e.g., Quine's Japanese classifier example (*DL*, pp. 35-38); but in the appendix I will suggest a device by which such apparent exceptions can be handled.

make physicalistic sense (call this "Requirement 2"). Many views about what partially refers to what that are compatible with Requirement 1 fail to meet Requirement 2. For instance, suppose someone were to hold that "*gavagai*" determinately signifies the set of rabbits (that is, partially signifies that set *and no other*), and that "*glub*" determinately signifies identity. This hypothesis would lead to reasonable conclusions about the truth and falsity of sentences; but if the Quinean argument sketched earlier in the paper is correct (that is, if we cannot make physicalistic sense of any semantic connection between "*gavagai*" and the set of rabbits that does not hold equally between "*gavagai*" and the set of undetached rabbit parts), then the hypothesis would be unsatisfactory since it would rule out the possibility of a physicalistic explanation of partial signification.

A similar point can be made for the notion of basis. Thus we could get reasonable results about truth and falsity by singling out one of the natural-kind terms of the language—for example, "*gavagai*"—and letting this term serve as the basis of all of the other natural-kind terms (and for the identity term). But this would be unreasonable since there seem to be no facts about the use of "*gavagai*" which give us objective grounds for singling it out for special treatment over all of the other natural-kind terms of the language. It seems then that if we are to avoid physically undetermined choices, we have to treat all natural-kind terms on a par; and the only way I know to do this and still get reasonable results about truth and falsity is to let all the natural-kind terms be dependent terms whose basis is "*glub*."[12]

I am suggesting, then, that we should decide what partially denotes what, what is the basis of what, and so forth, largely by

[12] Actually, there is a more positive reason for regarding the identity predicate as the basis for the other predicates (assuming we believe Quine's example). The reason is that the causal link alone is not sufficient to explain how "rabbit" could partially signify the set of rabbits and partially signify the set of undetached parts: perhaps the causal link could be used to explain how "rabbit" partially signified one or more sets whose union was {rabbits} ∪ {undetached rabbit parts}, but it is hard to see how the causal link could divide this union up into the two sets that are required. This problem does not arise if we view "rabbit" as acquiring its partial referents through identity considerations in the manner outlined on pp. 215-217.

the interaction of two considerations: considerations "from above" about what consequences those decisions have for truth and falsity, and considerations "from below" about what consequences the decisions have for the prospects of some day understanding semantics physicalistically. In other words, we should decide on the application of the semantic notions I have introduced in pretty much the same way that people have in the past decided upon the application of semantic notions like denotation and signification. The reason for replacing the notions of denotation and signification by the more general notions of partial denotation, partial signification, and so forth was that unless we made this generalization there would be situations in which it was impossible jointly to satisfy Requirement 1 and Requirement 2. (Thus before making the generalization we could satisfy only Requirement 1 by saying either that "*gavagai*" signified the set of rabbits or that it signified the set of undetached rabbit parts; while physicalistic considerations seemed to preclude "*gavagai*" from bearing any semantic relation to one of these sets without bearing it to the other set as well.) But I believe that *if we adhere to the semantics developed in this paper* (or the slight generalization of it suggested in note 11), there is no reason to doubt that Requirements 1 and 2 can always be jointly satisfied.

The semantic theory I have been developing is of some interest not only for Quine's radical indeterminacy thesis, but also for the philosophy of mathematics. Consider the following two ontological views:

(*i*) there are infinitely many physical objects, but no abstract objects;

(*ii*) there are sets in addition to physical objects, but no abstract objects other than sets.

On neither view are there any entities which are obviously numbers; therefore on both views there is a prima-facie problem of accounting for the truth and falsity of sentences expressed in number-theoretic terminology. But since both views admit the existence of infinitely many objects, the idea suggests itself of correlating an object with each numeral in such a way that no two numerals have the same correlated object. We could then say that each numeral denotes the object correlated with it, that a predicate such as "prime" signifies the set of objects correlated with "2," "3," "5," and so forth; and in this way we could get a definition of truth that did not require the existence of any objects except those allowed by (*i*) or by (*ii*).

In "What Numbers Could Not Be,"[13] Paul Benacerraf raised a difficulty for this reductivist proposal. The difficulty is that *many* correlations of the sort just discussed are possible, and it is clearly impossible to specify one correlation that gives "the real referents" of the number-theoretic words. The reductivist might try to escape this difficulty by saying that it is not important to his purposes to hold that number-theoretic terms referred to these correlated objects all along; it is sufficient (he might say) that we be able to *replace* number-theoretic talk by talk of the correlated objects. (Quine has suggested this line of escape in *Word and Object*, Sections 53 and 54.) But this will not do, for it suggests that earlier number theorists such as Euler and Gauss were not referring to anything when they used numerals, and (barring a novel account of mathematical truth) this would imply the unpalatable conclusion that virtually nothing that they said in number theory was true. Benacerraf's observation, then, seems to rule out the possibility of the reductivist proposal.

In fact, however, the reductivist has another answer to Benacerraf.[14] All that Benacerraf's observation shows is that if numerals are related to physical objects, they are not related to them in a determinate (that is, one-one) way. What the reductivist needs is an explication of number-theoretic truth in terms of correspondence relations between numerals on the one hand and physical objects and/or sets on the other; and a slight generalization of the semantics that work for Quine's "rabbit" example will work here as well.

The generalization is this: we need to be able to apply semantic locutions like "partially denotes" not just to individual terms, but

[13] *Philosophical Review*, LXXIV (1965), 47-73.

[14] Or more accurately, to philosophers who think that Benacerraf's observation shows that there are abstract objects distinct from sets—viz., numbers. In the final section of his paper Benacerraf himself expresses a reluctance to adopt such a platonistic position; so perhaps he would approve of the reductivist position suggested below.

to certain sequences of terms that can be naturally thought of as working together in a unit. Such sequences can be called *atomic sequences*, and we can stipulate that no term is contained in more than one atomic sequence. Then if we let the *atoms* consist of the atomic sequences and the terms that are not members of any atomic sequences, we can let the variable "t" in the (a) and (b) clauses of (7^*) range over atoms instead of over terms. The idea will become intuitively clear from the number-theory example. The first thing we want to say is that the ω-sequence of numerals *as a whole* is semantically related to certain ω-sequences of objects —namely, to those ω-sequences of which no two members are the same. Let us say, then, that the ω-sequence of numerals is an atomic sequence that partially denotes precisely the ω-sequences of objects just mentioned. Now what are we to say about number-theoretic predicates like "odd" and "prime"? The obvious answer is that they are dependent predicates whose basis is the sequence of numerals. *Relative to any ω-sequence σ* that is correlated with the numerals, the word "prime" signifies the set of objects in the prime positions of σ, and the word "odd" signifies the set of objects in the odd positions of δ. These seem to be rather natural stipulations, and they have the consequence that precisely the right number-theoretic statements come out true. Also, because we have adopted (8) instead of making arbitrary choices, the stipulations have the consequence that bizarre statements like "The number two is Julius Caesar" come out neither true nor false.

So the semantics of correlative indeterminacy shows that Benacerraf's observation is not ultimately of much relevance to the possibility of adhering to austere ontologies like (i) and (ii) and yet retaining the notion of mathematical truth. But my main concern in this paper is not with the value of the unrestricted truth-definition for the philosophy of mathematics, but with its value in showing that much of what Quine says about the significance of indeterminacy is wrong. My view about the significance of indeterminacy differs from Quine's view in two central respects. The first respect was stressed in the early pages of this paper: I view indeterminacy as showing not that a correspondence theory of truth is hopeless, but that a correspondence theory of truth must be a bit more complicated than we might have once thought; and I view indeterminacy as lending no support whatever to Quine's thesis of ontological relativity (which I have argued to be untenable). The second important respect in which my view differs from Quine's is that I do not see indeterminacy as showing any *arbitrariness* in semantic theory.[15] According to Quine, whenever we give a semantic theory for a foreign language, there is another semantic theory equally as good as the first, but incompatible with it. The reason Quine believes this is that on his view we can give a semantic theory for a language only after making an arbitrary choice among all the acceptable translation manuals mapping that language into ours; and if we had chosen a different manual, we would have gotten a different (and incompatible) semantic theory for the language. On my view, however, what we should do in this situation is to transcend the individual manuals: if many translation manuals are acceptable (or better, if many structures are relevant), then any semantic theory for the language that looks at only one of them is *inadequate*, and an *adequate* theory has to look at *all* of them. This has the consequence that now there are no alternatives to our semantic theory that are just as acceptable as it, but incompatible with it; therefore the objectivity of our semantic theory is restored.

Appendix

In note 5 I suggested that on my view, indeterminacy is a lot like vagueness: the semantics of partial denotation and partial signification provides an adequate treatment of both (except that for vagueness, it is necessary to quantify the semantics by talking of *degrees* of denotation and of signification, and hence degrees of

[15] I am excluding from consideration here the uninteresting kind of indeterminacy mentioned in the last paragraph of the appendix, the kind of indeterminacy which rests solely on the vagueness of current semantic terms; for it is true, but trivial, that any science whose terms are currently vague can be developed in more than one way equally well. The arbitrariness that Quine believes to exist is more radical, since he believes that arbitrary choices are necessary even apart from any vagueness that our semantic terms might possess. (Cf. last sentence of appendix.)

truth). It might be thought, however, that it is possible to treat vague predicates in the way I (following David Lewis) have suggested only if there are some predicates in our language which are not vague—namely, the predicates we use in specifying the partial extensions of the vague predicates. Analogously, it might be thought that it is possible to treat indeterminacy in our own language by the methods I have advocated only if there are determinate predicates in our own language to use in specifying the partial extensions of the indeterminate predicates. If this were so, then my treatment of Quine's example would be inadequate, for I tried to describe the partial extensions of the (allegedly) indeterminate term "rabbit" by using the (allegedly) indeterminate expressions "rabbit" and "undetached rabbit part."

The reason for thinking that we can use only determinate terms in describing the partial extensions of indeterminate terms is, presumably, that it is not clear how assertions like

(10) "rabbit" partially signifies the set of rabbits and partially signifies the set of undetached rabbit parts

could be understood by anyone who regarded the last two tokens of "rabbit" as indeterminate. But I will argue that though it is psychologically easier to understand assertions like (10) if we pretend that these assertions are phrased in a determinate meta-English, still they are just as intelligible, and have precisely the same truth conditions, if we remember that our meta-English is indeterminate.

To argue for this claim, it is necessary to know what a believer in indeterminacy would say about the expression "undetached rabbit parts." Relative to identity, of course, it signifies the set of undetached rabbit parts; but what does it signify relative to paridentity? Let us suppose that the answer to this question is that it signifies the set of *undetached parts of undetached rabbit parts*; and let us assume that undetached parts of undetached rabbit parts are *not* undetached parts of rabbits (or of anything other than undetached rabbit parts).[16] Let us now turn to a more

[16] Some such assumption is necessary if Quine's "*gavagai*" example is to work for a language that contains a predicate which is naturally translated

important question: what should a believer in indeterminacy say about the expression "partially signify"? Relative to identity, of course, it signifies the relation (R_1)

$\hat{x}\,\hat{y}$ (x partially signifies y).

But what does it signify relative to paridentity? To answer, note that "partially signify" is a two-place relation symbol whose left-hand argument ranges over abstract entities (namely, expression types) and whose right-hand argument ranges over sets of physical objects. For such expressions, we should expect the believer in indeterminacy to accept the schema

Relative to paridentity, "——" signifies the relation $\hat{x}\,\hat{y}$ (x ——'s the set of objects whose undetached parts are in y).

In particular, he will hold that relative to paridentity, "partially signifies" signifies the relation (R_2)

$\hat{x}\,\hat{y}$ (x partially signifies the set of objects whose undetached parts are in y).

It is clear then that if we evaluate (10) according to the semantics of correlative indeterminacy outlined in this paper, we will get that (10) is true if and only if

(*i*) "rabbit" partially signifies the set of rabbits and partially signifies the set of undetached rabbit parts,

and

(*ii*) "rabbit" partially signifies the set of objects whose undetached parts are in the set of undetached parts of rabbits, and partially signifies the set of objects whose undetached parts are in the set of undetached parts of undetached rabbit parts.

([*i*] comes from the structure in which "identical" is correlated with identity and R_1 is correlated with "partially signifies";

as "is an undetached part of." We can imagine the assumption true by stipulating that for something to be an undetached part of an x, it must be precisely half the size of an x.

[*ii*] comes from the structure in which "identical" is correlated with paridentity and R_2 is correlated with "partially signifies.") But (*ii*) is equivalent to (*i*), given the assumptions about the undetached-part-of relation made above: the weird treatment of "rabbit" and "undetached rabbit part" is canceled by the weird treatment of "partially signifies." We see then that (10) has the same truth conditions whether one pretends that meta-English is perfectly determinate or one recognizes that it too is indeterminate. I think it is clear that the same is true for the other sentences which I used in developing the semantics of indeterminate expressions, and for this reason I think that the objection to my approach fails.

I now want to consider a different objection to the semantics I have been advocating. This second objection, if it were valid, would also tell against Lewis-type treatments of vagueness; in fact, the objection is most easily formulated as an objection against such treatments of vagueness, so that is how I will formulate it here.

Consider sentences of the form

(S_h) All men of height greater than h are tall.

On a Lewis-type semantic theory there are two critical heights h_1 and h_2 (where $h_1 < h_2$), such that S_h is absolutely false if and only if $h \leqslant h_1$, and absolutely true if and only if $h \geqslant h_2$. The advantage of a Lewis-type treatment of vagueness over a standard semantic theory (one which does not recognize partial signification) is that Lewis's treatment accounts for the obvious fact that S_h does not jump *suddenly* from absolute falsehood to absolute truth as h increases; rather, there is an interval (between h_1 and h_2) during which S_h takes on successively higher degrees of truth *between* absolute falsehood and absolute truth. But it might be argued that although Lewis's treatment goes part way toward accounting for vagueness, it does not go far enough. For according to it, there is *a precise point* h_2 such that S_{h_2} is absolutely true while no sentences of the form $S_{h_2-\epsilon}$ are absolutely true (though some of them have extremely high degrees of truth). Suppose now that *another* Lewis-type theorist held that the critical point was really not h_2, it was h_2 — 0.001 inch; could we reasonably suppose that there was a fact of the matter as to which Lewis-type theorist was correct? Obviously not, and so it seems that Lewis-type theorists have merely put the problem of vagueness one step back: they explain why there is no fact of the matter as to whether $S_{\frac{1}{2}(h_1 + h_2)}$ is absolutely true or absolutely false, by saying that it is neither, but they do not explain why there is no fact of the matter as to whether $S_{h_2 - 0.001}$ is absolutely true.

If this objection to Lewis-type treatments of vagueness were valid, it seems likely that an analogous objection could be made to my treatment of indeterminacy (though the examples are considerably more complicated to describe). In fact, however, the objection is invalid. What the objection overlooks is that a Lewis-type theorist need not hold that "tall" is the only vague term that is relevant to the discussion; he can hold that the semantic terms "true" and "partially signifies" are also somewhat vague. The vagueness of "tall" serves various practical purposes; but the vagueness of our semantic terms serves no useful purposes, and as semantic theory develops we can expect these terms to become more precise. On *some* ways of making them precise, $S_{h_2 - 0.001}$ will be absolutely true; on *other* ways of making them precise, it will not be *absolutely* true but will have an extremely high degree of truth. (On *no* way of making them precise will $S_{\frac{1}{2}(h_1 + h_2)}$ be either absolutely true or absolutely false.) The fact that there is no fact of the matter as to whether $S_{h_2 - 0.001}$ is absolutely true, as we currently use "absolutely true," is simply due to the fact that our current use of "absolutely true" is imprecise.

This answers the objection to Lewis's semantics, and the analogous objections which might be raised against my own semantics. I should also mention that the idea utilized in answering the objection—the idea of positing vagueness in our semantic terms—is essential for dealing with some of Quine's less interesting examples of indeterminacy—for example, the Japanese classifier example. But there is no hope of treating the "*gavagai*" example as merely illustrating the vagueness of our semantic terms: for the difficulty that the "*gavagai*" example raised was that it seemed impossible to find any hypothesis about what "*gavagai*" signifies that satisfies both Requirement 1 and Requirement 2; and it is obvious that this difficulty (if it really exists) cannot be alleviated by making the term "signifies" more precise.[17]

HARTRY FIELD

Princeton University

[17] Adam Morton, Gilbert Harman, Paul Benacerraf, and the editors of the *Philosophical Review* made helpful comments on earlier versions of this paper. My work was supported by the National Endowment for the Humanities.