

Why Meaning (Probably) Isn't Conceptual Role

JERRY FODOR AND ERNEST LEPORE

Introduction

It's an achievement of the last couple of decades that people who work in linguistic semantics and people who work in the philosophy of language have arrived at a friendly, *de facto* agreement as to their respective job descriptions. The terms of this agreement are that the semanticists do the work and the philosophers do the worrying. The semanticists try to construct actual theories of meaning (or truth theories, or model theories, or whatever) for one or another kind of expression in one or another natural language; for example, they try to figure out how the temperature could be rising compatibly with the substitutivity of identicals. The philosophers, by contrast, keep an eye on the large, foundational issues, such as: what's the relation between sense and denotation; what's the relation between thought and language; whether translation is determinate; and whether life is like a fountain. Every now and then the philosophers and the semanticists are supposed to get together and compare notes on their respective progress. Or lack thereof.

Accordingly, this paper is about not semantics but the philosophy of language. Some of us have been poking around in the basement of meaning theory, and we seem to have discovered a large, nasty crack; as far as we can tell, one of the foundation stones is coming unstuck. We thought we'd better tell you about it before things get worse.

We'll proceed as follows: First we'll try to say where in the foundations the problem is located; then we'll try to say what the problem is; and then we'll make a suggestion or two about what to do about the problem. The first part of the discussion will be very broad; the second part will be rather more specific; the third part will be practically nonexistent.

Here goes.

1. Where the Problem Is

A traditional foundational problem in the theory of meaning is: Where do semantic properties come from? The presupposition of this question is that the fact that a word (or a sentence, or whatever) means what it does can't be a brute fact. It can't be a brute fact, for example, that 'dog' means dog and not proton and that 'proton' means proton and not dog. Rather, 'dog' must have some nonsemantic property in virtue of which it means dog and not proton; and 'proton' must have some (different) nonsemantic property in virtue of which it means proton and not dog. To put it in the standard philosophical jargon, semantic properties must supervene on nonsemantic properties. There may be some properties that things just have; that they have for no reason at all. But if there are, they are the kinds of properties that basic physics talks about (like mass, charge, and charm). They certainly don't include the kinds of properties that semanticists talk about (like meaning dog or being a synonym of 'bachelor').

We remark in passing that none of this is to be construed as an attempt to legislate physicalism. For present purposes we're content that semantic properties should be, for example, irreducibly intentional, or irreducibly epistemological, or irreducibly teleological. But we take it to be not on the cards that they are irreducibly semantic. In short, we don't care whether semantic properties supervene on something that is physical just as long as they supervene on something other than themselves.

So the question arises: What do the semantic properties of symbols supervene on? Over the years, philosophers of language have been enthusiastic for two quite different (perhaps, indeed, incompatible) kinds of answers to this question. There's what we will call the 'Old Testament' story, according to which the meaning of an expression supervenes on the expression's relation to things in the world; and there's what we will call the 'New Testament' story, according to which the meaning of an expression supervenes on the expression's role in a language. The disagreement between these two sorts of story is venerable, and we don't propose to go into the details here. Just a paragraph or two by way of assembling reminders.

Old Testament semantics derives most directly from the British Empiricists and has, among its modern representatives, behaviorist psychologists like Watson and Skinner and a handful of philosophers in the 'naturalized semantics' movement (including Dretske, Millikan, Barwise and Perry in

certain of their moods, and one or more of the Fodors). The basic idea is that 'dog' means *dog* because of some (nonsemantic) relation that holds between the symbol and the animal. If you are a British Empiricist (or, at least, if you're Hume), you say that this relation comes down to some variant on resemblance; 'dog' means *dog* rather than *proton* because 'dog' is associated to a certain mental image that means *dog* rather than *proton*. And the mental image means *dog* rather than *proton* because it resembles dogs quite a lot but resembles protons hardly at all. If you are a behaviorist psychologist, you say that the symbol-world relation that semantic properties supervene on is *causal* (typically associative); 'dog' means *dog* rather than *proton* because, in consequence of the speaker's history of conditioning (or something) dogs cause utterances of 'dog' and protons don't. Recent developments of Old Testament Semantics propose still other variations on this theme, including, for example, appeals to nomological and/or informational relations between symbols and the world.

Now, this Old Testament story about meaning has come under a lot of criticism, both in philosophy and in linguistics. Perhaps the basic objection is one that derives from Frege: Meaning *can't* be a symbol-world relation, according to this objection, because *identity* of symbol-world relations is compatible with *difference* of meaning. Thus, it's plausible that the expressions 'The Morning Star' and 'The Evening Star' are both attached to the same nonlinguistic thing (viz. to Venus); but it's also plausible that they don't *mean* the same thing. If they did, you couldn't deny that the Morning Star is the Evening Star without self contradiction or assert that it is without tautology. Both of which, in fact, you can.

So, then, according to this argument, the meaning of an expression doesn't supervene on the way that it's attached to the world. So what *does* it supervene on?

The New Testament story is an elaboration of the following idea: the expressions 'The Morning Star' and 'The Evening Star' mean different things, despite their both being attached to Venus, because they have different *roles in the (English) language*. 'Only in the context of a sentence does a word have a meaning' says Frege, and Wittgenstein adds that 'to understand a sentence is to understand a language'. The core idea is that it's the way they are connected *to one another* that determines what the expressions in a language mean. Very often (in fact, in all the semantic theories that we will discuss) this notion of 'role in the language' is given an epistemological twist: To master a linguistic expression is to be disposed to draw some core set of *inferences* that fix its semantic relation to other expressions in the language. To master the expression 'dog' is, *inter alia*, to be prepared to infer from 'Rover is a dog' to 'Rover is an animal'. To master the expression 'The Morning Star' is, *inter alia*, to be prepared to infer from 'is the Morning Star' to 'rises in the Morning'. And so forth.

We want to emphasize that, at this stage, we're using notions like *infer* and *disposed to infer* as blank checks. For the moment we'll let any disposition to have one belief cause another count as a disposition to infer the

second from the first. So deductive inferences, and inductive inferences, and plausible inferences, and prudential inferences, and mere associations, and Heaven knows what all else, are included *pro tem*. A lot of this paper will be about problems that arise when a New Testament Semanticist tries to say exactly which inferences are constitutive of the meanings of the terms that enter into them.

The New Testament picture in philosophy is quite close to one that's familiar from Structuralist linguistics, according to which the meaning of an expression is its role in a 'system of differences'. To know what 'dog' means is to know that it excludes 'cat' and cohabits with 'animal'. To know what 'bachelor' means is to know that it excludes 'spinster' and cohabits with 'unmarried.' Notice that it follows from this view that if English didn't contain words that mean what 'cat' and 'animal' do, *it couldn't contain a word that means what 'dog' does*. This is a kind of conclusion that structuralists in semantics have, with varying degrees of enthusiasm, quite generally been prepared to embrace.

Although, as far as we can tell, one or other version of the New Testament view is held practically as gospel not only in linguistics and philosophy, but throughout the Cognitive Sciences, it is nevertheless possessed of well known problems. For example, New Testament Semantics appears to be intractably holistic. This is because, once you start identifying the content of a belief with its inferential role in a system of beliefs (*mutatis mutandis*, the meaning of an expression with its inferential role in a language), it's hard to avoid proceeding to identify the content of a belief with its *whole* inferential role in the system of beliefs. And, having gone that far, it's hard to avoid the conclusion that if belief systems differ at all with respect to the propositions they endorse, then they differ completely with respect to the propositions they endorse. (*Mutatis mutandis*, if languages differ at all in respect of the propositions they can express, then they differ entirely in respect of the propositions they can express.) This is a well-greased, and well travelled, slippery slope; having arrived at the bottom, one finds oneself accepting such *prima facie* outlandish doctrines as that no two people ever share a belief; that there is no such relation as translation; that no two people ever mean the same thing by what they say; that no two time slices of *the same* person ever mean the same thing by what they say; that no one can ever change his mind; that no statements, or beliefs, can ever be contradicted (to say nothing of refuted); and so forth. It's a moot question how to get the goodness out of inferential role semantics without paying this extravagant price. Indeed, it's moot whether it is possible to do so.

Serious though these worries are, however, they are all, as it were,

' The issues we'll be concerned with arise in precisely parallel ways for theories of the meanings of linguistic expressions and theories of the contents of thoughts. In the rest of what follows, we will not bother to draw this distinction unless something turns on it.

external to semantics as such; they impugn its consequences for epistemology and ontology rather than its actual coherence. And a convinced New Testament Semanticist might be prepared to bite the bullet. If the doctrine that meaning supervenes on intralinguistic relations has relativistic, idealistic or even solipsistic consequences, then perhaps we had better learn to live with relativism, idealism or even solipsism. There are many, especially on the West Coast, who clearly long to do so.

Well, so much for the background; now here's the problem. It looks to us as though there are serious *linguistic* problems for the functional role approach to meaning. We want to argue that, quite aside from ontology and epistemology, it looks like New Testament semantics can't be squared with a pair of linguistic hypotheses that practically everybody thinks there are very good reasons to endorse; and which, *de facto*, most New Testament semanticists actually *do* endorse. These are, first, that natural languages are compositional, and, second, that the analytic/synthetic (a/s) distinction is unprincipled. (This means not that the distinction is *vague* - *what* empirical distinction isn't? - but that there aren't any expressions that are true or false solely in virtue of what they mean.) The foundational problem to which we wish to call your attention is that these two principles together with New Testament Semantics form an inconsistent triad. At least one of the three will have to go.

The next section of this paper sketches our argument for this claim; the third asks what on earth to do about it.

2. Compositionality, Analyticity and Inferential Role Semantics

There is, we believe, an internal connection between analyticity and compositionality that hasn't previously been remarked upon and that has dire implications for the New Testament identification of meanings with inferential roles. We commence our exploration of this territory with a few remarks about compositionality.

A language is *compositional* iff (idioms aside) the meaning of its syntactically complex expressions is a function of their syntactic structures together with the meanings of their syntactic constituents. For present purposes, a language is compositional iff the meaning of its sentences is a function of their syntactic structural descriptions together with the meanings of their lexical constituents. We take the doctrine that natural languages are compositional to be, as one says in Britain, non-negotiable.

No doubt connectionist psychologists (and a handful of philosophers) have occasionally been moved to dispense with compositionality. But that just indicates the desperateness of their plight. For, compositionality is at the heart of some of the most striking properties that natural languages exhibit. The most obvious of these are *productivity* (roughly, the fact that every natural language can express an open ended set of propositions) and *systematicity* (roughly, the fact that any natural language that can express

the proposition P will also be able to express many propositions that are semantically close to P. If, for example, a language can express the proposition that aRb, then it can express the proposition that bRa; if it can express the proposition that P → Q, then it can express the proposition that Q → P; and so forth.)

We digress to remark (since the issue often comes up) that pairs like 'John calculated the answer' and 'the answer calculated John' are *not* exceptions to the systematicity of English. Systematicity requires that, generally, if a language can express the proposition P and the proposition P is semantically close to the proposition Q, then the language can also express Q. If, however, there is *no* such proposition as Q, it is no objection to its systematicity that a language can't express it. We are inclined to think that there is no such proposition as that the answer calculated John, hence that the present examples do not constitute an objection to the systematicity of English. We do not, however, wish to dogmatize. If there is such a proposition as that the answer calculated John, then English *can* express it; indeed the form of words 'the answer calculated John' does so. English is thus systematic on either assumption.

Connected to both productivity and systematicity is a further, apparently perfectly universal, feature of natural languages. The structure of sentences is, in the following sense, *isomorphic* to the structure of the propositions they express: If a sentence *S* expresses the proposition that *P*, then syntactic constituents of *S* express the constituents of *P*. If for example, a sentence expresses the proposition that *P* and *Q*, then there will be one syntactic constituent of the sentence that expresses the proposition that *P* and another syntactic constituent that expresses the proposition that *Q*. If a sentence expresses the proposition that John loves Mary, then there will be a syntactic constituent of the sentence that refers to John, another syntactic constituent of the sentence that refers to Mary, and another syntactic constituent of the sentence that expresses a relation such that, necessarily, that relation holds between *x* and *y* iff *x* loves *y*. Notice that, though all of this is patent, none of it is truistic. Idioms and other 'holophrastic' constructions are all exceptions, albeit the sorts of exceptions that prove the rule.²

Our point is that these three generalizations about natural languages productivity, systematicity and isomorphism - are connected and explained on the assumption of compositionality; they are all consequences of the principle that the meaning of a sentence is composed from the

Our statement of the isomorphism principle is intentionally left pretty vague; deep issues arise when more precision is attempted. Suppose, for example, that you hold that (in a null discourse) the sentence 'it's raining' expresses the proposition that it's raining in the context of utterance. Then either you must say that 'its raining' has more constituents than appear on its surface or that the isomorphism principle can be violated by pragmatically carried information. For present purposes, we propose not to broach these sorts of issues.

meanings of its parts. And they appear to be completely baffling otherwise. So, we intend to insist on compositionality in what follows.

Now we get to the crack in the foundation. It turns out that compositionality is an embarrassment for the kind of New Testament semantics that identifies the meaning of an expression with its inferential role. In particular, it invites the following kind of prima facie argument:

- meanings are compositional
- but inferential roles are not compositional
- so meanings can't be inferential roles.

The second step is, of course, the one that's doing the work. But it seems pretty obviously sound. Consider the meaning of the phrase 'brown cow'; it depends on the meanings of 'brown' and 'cow' together with its syntax, just as compositionality requires. To a first approximation, 'brown' means - if you like, it connotes the property - BROWN, 'cow' means COW, and the semantic interpretation of the syntactic structure (ADJECTIVE + NOUN)N is property conjunction. (We are aware that there are problems about decoy ducks and the like; but the assumption that language is compositional is the assumption that such problems can be solved.) But now, prima facie, the inferential role of 'brown cow' depends not only on the inferential role of 'brown' and the inferential role of 'cow', but also on what you happen to believe about brown cows. So, unlike meaning, inferential role is, in the general case, not compositional.

Suppose, for example, you happen to think that brown cows are dangerous; then it's part of the inferential role of 'brown cow' in your dialect that it does (or can) figure in inferences like 'brown cow - dangerous'. But, first blush anyhow, this fact about the inferential role of 'brown cow' doesn't seem to derive from corresponding facts about the inferential roles of its constituents. You can see this by contrasting the present case with, for example, inferences like 'brown cow --* brown animal' or 'brown cow --> not-green cow'. 'Brown animal' follows from 'brown cow' because 'cow' entails 'animal'; 'not-green cow' follows from 'brown cow' because 'brown' entails 'not-green'. But it doesn't look like either 'brown' or 'cow' entails 'dangerous,' so, to this extent, it doesn't look like the inference from 'brown cow' to 'dangerous' is compositional.

In short, it appears that some, but not all, of the inferential potential of 'brown cow' (some of its role in one's language or belief system), is determined by the respective inferential potentials of 'brown' and 'cow', the rest being determined by one's 'real world' beliefs about brown cows. This should not seem surprising or contentious; it's just a way of saying that 'brown cows are dangerous' (unlike 'brown cows are animals' and 'brown cows are not-green' and 'brown cows are brown') is clearly synthetic. That is 'brown cows are dangerous' is contingently true, true in virtue of the facts about brown cows, not in virtue of the facts about meanings (assuming that it's true at all).

But, to repeat, if meanings are compositional and inferential roles aren't, then it follows that meanings can't be inferential roles.

As far as we can tell, this line of argument is quite robust; in particular, it doesn't depend on detailed assumptions about how an inferential role semantics construes the notion of inferential role. There is, for example, an influential paper by Hartry Field in which inferential role is analyzed in terms of subjective probabilities; in effect, the inferential role of your thought that P is identified with the subjective probability that you (would) assign to P contingent on each of the other thoughts that you (can) entertain (Field, 1977). So, for example, the inferential role of your thought that it's raining is determined in part by the subjective probability that you (would) assign to that thought on the assumption that the streets are wet; and in part by the subjective probability that you (would) assign to it on the assumption that the sun is out; and in part by the subjective probability that you (would) assign to it on the assumption that elephants have wings . . . and so forth. Like other species of New Testament Semantics, this treatment generalizes, in fairly obvious ways, from a theory according to which inferential roles are assigned to thoughts to one in which they are assigned to linguistic expressions, or to both.

Our point is that the construal of inferential roles in terms of subjective probabilities, whatever other virtues it may have, does nothing to help with the compositionality problem. This is because *subjective probabilities are not themselves compositional*. For example, the subjective probability one assigns to the thought that *brown cows are dangerous* is not a function of the subjective probability one assigns to the thought that *cows are dangerous*, together with the subjective probability that one assigns to the thought that *brown things are dangerous*. If this seems not obvious, consider a world (or rather a belief world, since the probabilities at issue are supposed to be subjective) where there are very many things that are cows, almost none of which is dangerous, and very many things that are brown, almost none of which is dangerous, and a very small number of brown cows, almost all of which are very, very fierce. On these assumptions, the probability that something that is brown is dangerous is small, and the probability that something that is a cow is dangerous is small, but the probability that a brown cow is dangerous is as big as you please.

We intend the argument so far as one horn of a dilemma, and we anticipate the following reply:

OK, so if the compositionality of meaning is assumed, meanings can't be identified with inferential roles as such. But this doesn't *really* embarrass New Testament Semantics because meanings can still be identified with roles in *analytic* inferences. Thus, on the one hand, the inference 'brown cow ----> brown animal' is compositional (it's inherited from the inference 'cow -> animal') and, on the other hand, precisely because it is compositional, 'brown cow -j brown animal' is analytic. Compositional inferences will always

be analytic and analytic inferences will always be compositional; the compositionality of an inference and its analyticity *are the same thing*.

Look at it this way: If 'brown cow -> brown animal' is compositional, then it's warranted by the inferential roles of the expressions 'brown' and 'cow'. That's what it is for an inference to be compositional. But, according to New Testament Semantics, the inferential roles of 'brown' and 'cow' *are their meanings*. So, then, that 'brown cow -> brown animal' is warranted follows from the *meanings* of 'brown' and 'cow'. But for an inference to be analytic *just is* for it to be warranted by the meanings of its constituent expressions. So the compositionality of 'brown cow -> brown animal' - or, *mutatis mutandis*, of any other inference -*entails* its analyticity. The same argument also works the other way around. For an inference to be analytic is for it to be warranted by the meanings of its constituents. But, according to New Testament Semantics, meanings are inferential roles. So, for an inference to be analytic is for it to be warranted by the inferential roles of its constituents. But for the warrant of an inference to be determined by the inferential roles of its constituents is for the inference to be compositional. So compositionality entails analyticity and vice versa. So, then, meaning is compositional and inferential role isn't and role in analytic inference is. What all that shows is just that we need a revised version of New Testament Semantics; one which identifies meaning with role in analytic inference.

The first thing to say about this new suggestion is that the threat of circularity is now very close to the surface. It is proposed that we reconcile New Testament Semantics with the compositionality of meaning by identifying the meaning of an expression with its role in analytic inferences. But the difference between analytic inferences and inferences *tout court* is just that the validity of the former is guaranteed *by the meanings* of their constituent expressions. So analyticity, meaning (and compositionality) scrape out a living by doing one another's wash, and Quine gets to say 'I told you so'.

Notice also that the naturalizability of inferential role semantics is jeopardized by the present proposal. A lot of the attraction of identifying meaning with inferential role lies in the thought that the inferential role of an expression might in turn be identified with causal role, thereby conceivably providing the basis for a naturalistic solution to Brentano's problem. That causal relations reconstruct inferential relations is a foundational assumption of computational theories of mental processes, so perhaps there is hope here of a unification of semantics with psychology. But, barring proposals for a causal theory of analyticity, this tactic is unavailable to the philosopher who identifies meaning with the role of

an expression in analytic inference.' The idea that mental processes are computational may provide the basis for a naturalistic account of inference, but it offers no insight at all into the nature of analyticity. Nor, as far as we can see, does anything else.

We can now say pretty exactly what our problem is: You can't identify meanings with inferential roles tout court, since, unlike meanings, inferential roles tout court aren't compositional. You can identify meanings with roles in analytic inferences, however, because analytic inferences are compositional. But, of course, the cost of identifying meanings with roles in analytic inferences is buying into the analytic/synthetic distinction. So the cost of New Testament Semantics is buying into the analytic/synthetic distinction. But, these days, practically everybody thinks that the analytic/synthetic distinction is unprincipled. Indeed, it's widely thought that the discovery that the analytic/synthetic distinction is unprincipled is one of the two most important achievements of modern philosophy of language; the other being precisely the theory that meaning supervenes on inferential role. If, then, we continue to assume that compositionality is non-negotiable, it follows that one of these foundational principles of the philosophy of language is going to have to go. Which one? And with what are we to replace it?

3.

What Now?

This is where we ought to tell you how to go about mending the cracked foundations; and we promise you that we would if we could. Since we can't, however, we'll restrict ourselves to a few more or less jaundiced remarks on what we take the visible options to be and on what we think that adopting one or other of them is likely to cost.

The first possibility one might consider is to try resuscitating the analytic/synthetic distinction. There is, we think, a certain face plausibility to this suggestion. We've been seeing that compositionality and analyticity come to much the same thing so long as you accept the New Testament view that meaning supervenes on inferential role. But, in fact, it's plausible that compositionality entails analyticity whether or not you accept New Testament Semantics. So, for example, it's hard to see how anybody could claim that the meaning of 'brown cow' is compositional while denying that the inference from 'brown cow' to 'brown' is analytically valid. If it is undeniable that the meaning of 'brown cow' is constructed from the

There are causal theories of meaning around; see, for example, Skinner (1957), Dretske (1981) and Fodor (1991) among many others. And each of them implies a corresponding notion of analyticity. But all these theories are externalist and atomist and thus offer no comfort either to New Testament Semantics or to holism. They don't legitimate a construal of meaning in terms of analytic inference because they reject New Testament Semantics; they don't reconstruct meaning in terms of inference at all.

meanings of 'brown' and 'cow,' it seems equally undeniable, and for the same reasons, that the inference from 'brown cow' to 'brown' is guaranteed by the linguistic principles that effect this construction. But an inference that is guaranteed by linguistic principles *just is* an analytic inference. In short, the very structural relations among the constituents of a sentence that ground its compositionality would appear to engender the analyticity of some of the inferences in which its constituents are involved.

So, to repeat, it appears that compositionality underwrites certain analyticities all by itself, without further appeal to the principle that meaning supervenes on inferential roles. A New Testament Semanticist might thus reasonably argue that if compositionality is non-negotiable, so too is analyticity; hence that the rejection of Quine's anti-analyticity arguments is independently motivated. (It's worth mentioning in this respect, that Quine says little or nothing about examples like 'brown cow ---> brown', where the analyticity of an inference depends on the *structure* of its premises rather than their lexical contents.) But if the revival of the analytic/synthetic distinction can be defended, it looks as though everything is all right. We've seen that you can't have New Testament Semantics and compositionality *and the rejection of als*, but you certainly can have the first two if you drop the third.

We emphasize, however, that the kind of a/s distinction that compositionality underwrites holds only between expressions and their *syntactic constituents*. It thus serves to distinguish, say, 'brown cow ---> brown' from 'brown cow --j dangerous.' But it does *not* underwrite a distinction between, say, 'brown cow --> animal', and 'brown cow --> dangerous.' That is, it doesn't underwrite an analytic/synthetic distinction among inferences that turn on the lexical inventory of the premises as opposed to their linguistic structure. It is, however, precisely these lexically governed inferences that make trouble for New Testament Semantics.

If the meaning of 'brown cow' derives from the meanings of 'brown' and 'cow' (as it must if the meaning of 'brown cow' is compositional) and if the meaning of 'cow' is its inferential role (as it must be if New Testament Semantics is right), then it must somehow be arranged that 'brown cow' inherits the inference 'animal' *but not the inference 'dangerous'* from the meanings of its constituents. But that requires that we exclude from the semantic representation of 'cow' such information as, for example, 'cow --> kind of x such that the brown x-s are dangerous' (and analogously, we must exclude from the semantic representation of 'brown' such information as 'brown --- x such the cow x-s are dangerous'). But to say that we must treat these inferences as excluded from the semantic representations of 'cow' and 'brown' is just to say that they are *not* constitutive of the meanings of 'cow' and 'brown'. But to say *that is* to presuppose an analytic/synthetic distinction that applies to lexically governed inferences. And it is precisely an analytic/synthetic distinction for lexically governed inferences that the familiar Quinean polemics are widely supposed to jeopardize.

We've been putting a lot of weight on the difference between analyticities that are engendered by the compositional structure of an expression and those that are generated by the meanings of items in its vocabulary. This distinction is, of course, undermined if you assume that there is a level of representation at which lexical items are semantically decomposed: to posit such a level is, in effect, to hold that lexical meaning is itself compositional. So if there is lexical decomposition, then we can identify the meaning of an expression with its role in those inferences that are determined by its compositional structure *including the compositional structure of its lexical items*. The effect is to assimilate inferences like 'brown cow -)- animal' to inferences like 'brown cow ---> brown' since, at the level of semantic representation, both involve relations between expressions and their constituents. And both are distinguished from 'brown cow ---> dangerous' since, presumably 'dangerous' isn't a constituent of the representations of 'brown cow' at *any* linguistic level.

But, of course, to resolve the present worries by taking lexical decomposition for granted simply begs the question against Quine. If lexical meanings were compositional, then lexical items would be definable; and it's part of Quine's story about there being no principled a/s distinction that there is no principled notion of definition. If Quine's arguments show anything, they show that there is no way to reconstruct the intuition that 'brown cow -* animal' is definitional and 'brown cow --> dangerous' isn't.

We conclude that, although there may be reasons for resuscitating the analytic/synthetic distinction, the non-negotiability of compositionality isn't one of them. Compositionality licenses a distinction between 'brown cow -j brown' and 'brown cow --> dangerous' but not between 'brown cow dangerous' and 'brown cow -* animal'. So the original problem stands: if compositionality isn't negotiable, then either there is an analytic/synthetic distinction for lexically governed inferences (contra Quine), or inferential roles aren't what meaning supervenes on (contra New Testament Semantics).

Our second point is that the present situation is rife with ironies. We remarked above that once you say that the meaning of an expression supervenes on its inferential role, it's hard to stop short of saying that the meaning of an expression supervenes on its *whole* inferential role. That is, it's hard to stop short of relativizing the meaning of an expression to the whole language that contains it, with the consequences that expressions in different languages are semantically incommensurable. There is a sort of linguistic idealist who delights in these holistic implications of New Testament Semantics; you see more than a trace of this sensibility among the 'cognitive linguistics' crowd, among connectionists, and, of course, among philosophers like Rorty, Putnam, Kuhn and Derrida. It now appears, however, that inferential role semantics *doesn't have* any holistic implications after all, so relativism loses however the argument turns out.

Here, roughly, is how the argument from inferential role semantics to semantic holism was supposed to go:

Premise 1, New Testament Semantics: The meaning of an expression is at least partially constituted by the expression's inferential relations.

Premise 2, No A/S Distinction: There is no principled distinction between those of its inferential relations that constitute the meaning of an expression, and those that don't.

Conclusion, Semantic Holism: The meaning of an expression is constituted by all of its inferential relations, hence by all of its role in a language.

There are, we think, lots of reasons to disapprove of this way of defending semantic holism; not least that it depends, apparently ineliminably, upon a form of slippery slope argument, and these are notorious for leading from true premises to false conclusions. (That is, it depends on arguing from 'there's no principled difference between the Fs that are Gs and the Fs that aren't' to 'either none of the Fs are Gs or all of them are'.) For present purposes, however, we're prepared to ignore all that. The point we want to emphasize is that the argument can't be better than its second premise; that is, the argument from inferential role semantics to semantic holism explicitly depends on denying a/s.

So, the situation is this: if the analytic/synthetic distinction is principled, then the second premise is false and the argument for semantic holism is unsound. But if the analytic/synthetic distinction is unprincipled, then the first premise is false. This is because, as we've been seeing, respecting compositionality requires that the meaning of an expression be identified with its role in analytic inferences; and compositionality isn't negotiable. But the argument that there is no analytic/synthetic distinction is the argument that there aren't any analytic inferences. We assume the law of excluded middle - either there is an analytic/synthetic distinction or there isn't. Either way, the argument from inferential role semantics to semantic holism has to be unsound. We suppose that this demonstration should darken the skies of West Coast semantic holists. So be it; they need the rain.

The foundational problem, to recapitulate, is that you can't do all of the following: endorse inferential role semantics AND endorse compositionality AND abandon the a/s distinction. If compositionality is non-negotiable, that leaves you with only two options: endorse the a/s distinction or reject the idea that meaning supervenes on inferential role. We can't tell you which of these it would be best to do, but we do want to insist that wriggling isn't likely to get you off the hook. In particular, lots of cognitive scientists have hoped to reconcile themselves to the Quinean arguments by opting for a graded, or a contextualized, or otherwise denatured notion of analyticity. The idea is that, although Quine may have shown that the notion of identity of meaning is in trouble, it's still open that a notion of similarity of meaning might be evolved. It's supposed that grading the analytic/synthetic distinction would allow us to have both compositionality and the supervenience of meaning on inferential role. New Testament semantics would thus be vindicated, albeit in a sort of soft-edged version.

Nobody has shown that this can't be done; but nor does anybody have the slightest idea of how to do it. It seems to us that the arguments against a principled notion of identity of meaning work just as well against a principled notion of similarity of meaning, but we won't try to make that case here.⁴ Suffice it for present purposes to remark that the connections between analyticity and compositionality that we've been examining make the search for a graded notion of analyticity look unpromising. Compositionality is, after all, a principle that governs relations *among meanings* (it governs the relation between the meaning of a complex expression and the meanings of its constituents). So, if your semantic theory reconstructs meaning in terms of analytic inference, and if you have only a graded notion of analyticity, *then you will have to live with a graded notion of compositionality as well*. But what would a graded notion of compositionality be like? And, in particular, how would such a notion do what compositionality is required to do; viz. account for systematicity, isomorphism and productivity?

Wouldn't a graded notion of compositionality entail, at best, that a finite acquaintance with a language is adequate to *sort of understand* expressions not previously encountered? Or that if a language is capable of expressing the proposition that aRb, then it is *sort of* capable of expressing the proposition that bRa? Or that if the sentence S expresses the proposition P, then the constituents of S *sort of* express the constituents of P? But is there any sense to be made of such claims as, for example, that

('John loves Mary' sort of expresses the proposition that John loves Mary) only if 'John' sort of refers to John?

These are muddy waters and we do not recommend that you wade in them.

Here's where we've gotten to: Compositionality is non-negotiable; inferential roles are compositional only if the analytic/synthetic distinction is tenable for inferences that are governed by the lexical inventory. So, if Quine was right about a/s, the only remaining alternative is to give up on the New Testament idea that the meaning of lexical items is constituted by their inferential roles. Barring some proposal for an entirely new kind of semantics, this requires going back to the Old Testament view that what makes 'dog' mean *dog* is some sort of symbol-world connection; perhaps some sort of causal or informational or nomological connection between tokens of the expression type and tokens of the animal type. There are, as remarked above, various accounts of this currently on offer and, who knows, maybe one of them can be made to work.

The spinoffs for epistemology would be of some interest since the plausible candidates for the semantically relevant symbol to world relations

⁴See Chapters 1 and 7 of Fodor and Lepore (1992).

all look to be *atomistic*. If, for example, 'dog' means *dog* because dogs cause 'dog's, then it looks like your language could have a word that means what 'dog' means in English even if your language didn't have any other words that mean what English words do. *Prima facie*, in fact, it looks as though your language could have a word that means what 'dog' means in English even if it doesn't have any other words *at all*. All sorts of holistic arguments against the possibility of translation (and, more generally, against the commensurability of languages and ideologies) would then appear to be likely to collapse. More dark skies on the West Coast.

We aren't, to repeat the point one last time, actually urging you to follow this course. On the contrary, a lot of linguists will have good reason for not wanting to do so. For example, if Old Testament semantics is essentially right, what becomes of such linguistic studies as 'lexical semantics'? You can't study the semantic relations among lexical items unless there *are* semantic relations among lexical items; and it's the burden of Old Testament semantics that there aren't. Old Testament semantics has it that semantic relations hold between lexical items *and the world* and *only* between lexical items and the world.' Old Testament semantics has, therefore, no truck with semantic fields, lexical decompositions, conceptual networks and the like. If, indeed, what you mean by 'the semantic level of linguistic description' is the level of description at which sentences that differ only in synonymous expressions are identically represented, the natural way to read the Old Testament is as denying that there is a semantic level of linguistic description. So generative semantics goes, and interpretive semantics *also* goes, and the highest level of linguistic description is, as it might be, syntax or logical form; viz. a level where the surface inventory of nonlogico-syntactic vocabulary is preserved. The God of the Old Testament is an austere God.

As previously remarked, our present purpose isn't to sell you any such chilly theology; just to point out that there seems to be this crack in the foundations of the structure in which semanticists and philosophers of language have recently been cohabiting. And to urge that somebody do something about it not later than sometime before the roof falls in.

More precisely, the Old Testament view holds that, if there are semantic relations among lexical items, they are derived not from relations between their roles in the language but from relations between their connections to the world; *ceteris paribus*,

References

- Dretske, F. 1981: *Knowledge and the Flow of Information*. Cambridge, MA.: MIT Press.
- Field, H. 1977: Logic, Meaning, and Conceptual Role. *Journal of Philosophy*, 74, 379-409.
- Fodor, J. 1991: *A Theory of Content and Other Essays*. Cambridge, MA.: MIT Press.
- Fodor, J. and Lepore, E. 1992: *Holism, A Shopper's Guide*. Oxford: Blackwell.
- Skinner, B. F. 1957: *Verbal Behavior*. New York: Appleton Century Cross Inc.

two expressions connected to the world in the same way will be synonyms. But it is, of course, precisely the inferential relations among lexical items, and not their symbol world relations, that the semantic level of lexical description has always been supposed to represent.