

Article

Does an Inferential Role Semantics Rest Upon a Mistake?

PAUL A. BOGHOSSIAN

1. Introduction

Let's suppose that we think in a language of thought and that there are causal facts of the following form: the appearance in O's belief box of a sentence S_1 has a tendency to cause the appearance therein of a sentence S_2 but not S_3 . Ignoring many complications, we may describe this sort of fact as consisting in O's disposition to *infer* from S_1 to S_2 , but not to S_3 .

Let's call the totality of the inferences to which a sentence is capable of contributing, its *total inferential role*. A sub-sentential constituent's total inferential role can then be defined accordingly, as consisting in the contribution it makes to the total inferential role of the sentences in which it appears.

Against this rough and ready background, an *inferential role semantics* is just the view that there is some construct out of an expression's total inferential role that constitutes its meaning what it does. Let us call this construct an expression's *meaning-constituting inferential role*, or MIR, for short.¹

I am grateful to Ned Block, Jennifer Church, Stephen Schiffer and the members of my seminar on Analyticity at the University of Michigan, for helpful discussion and/or comment. Address for correspondence: Department of Philosophy, New York University, 100 Washington Square East, New York, NY 10003, USA.

Email: boghossn@acfcluster.nyu.edu (Internet) or boghossn@nyuacf (Bitnet).

¹ Although Fodor and LePore's title suggests that their paper is to be about *conceptual* role theories in general, rather than just about *inferential* role theories—inferential role being the purely linguistic aspect of an expression's overall causal role, a role that will typically include relations to non-linguistic stimuli or objects—their argument concentrates exclusively on the inferential part. For the sake of simplicity, therefore, I have framed the discussion as concerning inferential role theories only.

As no one will care to deny, at the present time an inferential role semantics (hereafter, IRS) is no more than a twinkle in the eyes of its advocates. To this day, no one has proposed a version that is detailed enough to be called a 'theory'; in particular, no one has even come close to providing a plausible specification of the constructs that are to constitute MIRs.

2. Fodor and Lepore's Argument

According to Fodor and Lepore, however, no one should even bother trying. For F & L claim to have on hand a powerful argument to the effect that *any* attempt to specify an expression's MIR will founder on one of two insuperable difficulties. The crux of their argument is neatly summarized in the following passage (1991, p. 337):

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 But, these days, practically everybody thinks that the analytic/synthetic distinction is unprincipled.

The *structure* of their argument seems pretty clear (though, as we shall see, its *content* is somewhat less transparent). An inferential role semantacist has to specify what is to constitute the MIR of an expression. And here he would appear to have only two choices: to identify the MIR of an expression with its total inferential role (what the authors refer to as its 'inferential role *tout court*'); or to identify its MIR with some proper subset of that total role. However, the authors contend, each of these options faces decisive objections.

Identifying the MIR of an expression with its total role leads to a theory of meaning that violates compositionality; and that is unacceptable. That leaves open the option of identifying it with some proper subset of that total role. But which subset should that be? The only plausible suggestion—indeed, the only suggestion that respects compositionality—is the identification of an expression's MIR with its role in *analytic* inference. However, and unfortunately, 'these days practically everybody thinks that the analytic/synthetic distinction is unprincipled'.

And so, the authors conclude, there is probably no satisfactory way to implement an inferential role semantics. Perhaps it is time that the philosophy of language reassessed the plausibility of an old-fashioned world-word semantics.

Now, I am not going to discuss F & L's argument for the first horn of their dilemma. I do not believe that a holistic version of an IRS—the version that identifies the MIR of an expression with its *total* role—violates compositionality. But since my reasons here mostly recapitulate Ned Block's, I will simply refer the reader to his contribution to this symposium.

However, given the many *other* problems with which a holistic version is beset—not the least being the well-known problems concerning change of mind, translation, and subsumption under intentional laws—considerable interest attaches to F & L's argument for their second horn. In my view—and *pace* Ned Block—the only plausible version of an inferential role semantics is a non-holistic version. It is important, therefore, to assess the effectiveness of F & L's argument against that option.

What, precisely, is that argument? As we saw, it consists in essence of the following:

1. The only plausible version of a non-holistic IRS identifies an expression's MIR with its role in analytic inference.
2. The analytic/synthetic distinction is 'unprincipled'.

Therefore

3. There can be no plausible version of a non-holistic IRS.

3. *The First Premise*

Let's begin by looking at the first premise. There are in fact *two* ways of understanding (1), depending on how the phrase 'identifies an expression's MIR with its role in analytic inference' is understood. On the one hand, the phrase might simply mean: *identifies an expression's MIR with its role in analytic inference, under that very description*. On the other hand, it might bear the following more complicated interpretation: *identifies an expression's MIR with its role in inferences which possess property P, where P is necessarily equivalent to the property of being analytic*. It is important, in the present connection, to distinguish between these two readings, and the distinct versions of an IRS that they determine.

Now, it is one of the more curious aspects of F & L's paper that they seem to intend that (1) be understood under its *simple*, rather than its complicated, interpretation. That is, the version of a non-holistic inferential role semantics that they take to be most plausible, and against which their argument is primarily directed, is a version that identifies an expression's MIR with its role in analytic inference, *under that very description*.

That this is the way they are construing the view under discussion is made very clear by what they have to say about it. Thus (p. 336):

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.

Furthermore (pp. 336–7):

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. . . . 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.

As these passages make amply clear, F & L intend their argument to be directed against an IRS that proposes to specify (some expression) e's MIR by saying that it consists in: the set of e-involving inferences whose validity is secured solely by the meanings of the participating expressions. No other hypothesis could explain their remarks. *Only* an IRS that employs a *non-reductive* notion of analyticity in its specification of MIRs faces a 'threat of circularity' or a 'fear that the naturalizability of semantics is [thereby] jeopardized'. And that is the only version of an IRS that they describe.

The trouble, however, is that I don't see how a viable IRS could *conceivably* assume the form that F & L describe. For I would have thought that with the sort of view they have in mind the 'threat of circularity' is not merely 'very close to the surface' but right smack on top of it for all to see; and that, as such, it simply is not a serious option. In other words, it seems to me that regardless of the ultimate cogency of the description 'analytic inference', it just *isn't* open to an IRS theorist to specify an expression's MIR by the mere use of that very description—by simply saying that an expression's meaning is determined by its role in inferences that are valid by virtue of its meaning. If an IRS theorist wishes to identify an expression's MIR with its role in analytic inference, he must identify it with its role in analytic inference *reductively* construed.

F & L may have been misled into thinking otherwise by their insistence that, for the purpose of the present discussion, they do not wish to presuppose 'physicalism'. Thus (p. 329):

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.

Putting aside objections to the slide from talk of reduction to talk of supervenience, the trouble is that any IRS that offers nothing more than an unreduced notion of analyticity in its definition of MIRs—as F & L's targeted IRS proposes to do—would not have provided *any* account of what, apart from the semantic, semantic properties are supposed to be identical with or supervenient upon. It's one thing to say that an IRS is entitled to use an unreduced notion of *inference*; it's quite another to say that it is entitled to use an unreduced notion of *validity by virtue of meaning*. The former doesn't necessarily threaten the informativeness of the relation it is supposedly trying to specify; the latter most certainly does.

I am inclined to think, therefore, that there is a serious question whether theoretical space even *allows* for the sort of view that F & L designate their most worthy opposition. For present purposes, however, I intend to waive this point. After all, a reductive version of an analyticity-based IRS survives it; and the soundness of F & L's second premise—and hence of the anti-IRS argument that they base on it—doesn't turn on whether the implicated notion of analyticity is construed reductively or not.

What, then, about the substance of their claim that the 'most plausible' version of a non-holistic IRS would identify e's MIR with the set of analytic inferences involving it? For reasons that will emerge below (in my discussion of the *Plausible-Sounding Principle*), I do not believe that F & L provide any support for this claim. But here again I propose to be concessive and to grant it to them. My worries about their argument are concentrated on their second premise, rather than on the first.

4. The Second Premise

So let us turn to looking at that second premise. To begin with, what does it mean to say that the analytic/synthetic distinction is 'unprincipled'? The authors provide the following explicit explanation:

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 by virtue of what they mean. (p. 332)

... the argument that there is no analytic/synthetic distinction is the argument that there *aren't any* analytic inferences.
(p. 340; emphasis in original)

Notice that according to the explanation provided, the notion of an 'analytic inference' is a purely *semantical* idea: to say that an inference is analytic is just to say that facts about its meaning suffice to fix its validity. In particular, the notion as defined appears to contain no *epistemological* component, mandating that analytic inferences be knowable in a certain way, for instance *a priori*.

F & L's idea, then, seems to be this: Quine showed (or at a minimum gave very good reason to believe) that there aren't any inferences (/sentences) that are valid (/true) solely by virtue of what they mean; and this is inconsistent with the claim that the meaning of an expression is determined by its role in analytic inference.²

Well, I am prepared to concede the second part of this claim, and hence the *validity* of the argument outlined. If it were true that the extension of the predicate 'analytic inference' were empty, then the IRS theorist targeted by F & L *would* be in serious trouble. Obviously. For this theorist seeks to explicate *e*'s meaning in terms of its meaning-constituting inferential role; and he seeks to specify that role by saying that it consists in the set of *e*-involving analytic inferences (construed either reductively or non-reductively). If, however, there are no analytic inferences, then every expression gets assigned the same 'null' meaning, and that is clearly unacceptable.

The only snag with this argument is that I don't think it can plausibly be claimed that we have been given—by Quine or by anyone else—good reason to believe that there are no analytic inferences. At a minimum—and more importantly for present purposes—F & L can't claim that we have.

To see why, we need to look briefly at a train of thought that I have developed at greater length elsewhere (Boghossian, forthcoming). Let us suppose—what is surely common ground between the principal disputants in *this* debate—that meaning realism is true and hence that there are determinate facts about what means what. Then, condensing considerably, there are bound to be facts about whether two expressions do or do not mean the same. Furthermore, if we simplify and suppose that an inference is analytic just in case it is transformable into a logical truth by the substitution of synonyms for synonyms (leaving aside the case where it is analytic as a result of a unidirectional conceptual entailment), then it follows that there are bound to be facts about whether the inference from

² Notice that, as I said above, this objection doesn't depend on whether the inferential role semantics is thought of as leaning on a reduced or unreduced notion of analytic inference. If there simply aren't *any* analytic inferences, then there aren't any reduced analytic inferences either.

This is e

to

This is e'

is or is not analytic (valid by virtue of meaning alone).³

At a minimum then, the meaning-realist must suppose that 'analytic inference' expresses a coherent, determinate property in good standing. Moreover, it is hard to see how he could plausibly avoid believing that it has a non-empty extension, that there actually are inferences that instantiate the property of being valid by virtue of meaning alone. For it is hard to see how it could plausibly turn out that, although the property of meaning identity is coherent and determinate, no two expressions could, either in point of principle, or in point of fact, instantiate it. (For the arguments, again see Boghossian, forthcoming.)

Making a long story very short, then, I don't see that any meaning-realist can accept the claim that there are no inferences valid by virtue of meaning alone. *A fortiori*, I don't see that F & L can.

5. Block's Objection

Ned Block has objected to this train of thought (see his contribution to this symposium, this issue, pp. 1–26). He accepts it, he says, for *wide* meanings, but not for *narrow* meanings. Why? According to Block, the point is (this issue, pp. 18, 20):

that narrow contents are *never* analytic. They aren't even the *sort* of thing that can be analytic. To see this, note that narrow contents are never *true*, and hence have no truth conditions. Analyticity requires truth—truth in all conditions, so something that can't be true can't be analytic either. . . .

Narrow meanings don't determine truth conditions at all, so Boghossian's argument is stopped short—at least for narrow meaning.

The point, if true, would be particularly *à propos* in the present context

³ Throughout this paper, the phrases 'valid by virtue of meaning alone' and 'true by virtue of meaning alone' are to be understood as meaning valid (or true) by virtue of meaning *and* logic alone. In 'Analyticity' (Boghossian, forthcoming) I distinguish between this 'impure' idea of analyticity and the 'pure' version that construes even the truths of logic as non-trivially true by virtue of meaning alone. This latter *positivist* notion of analyticity is not at issue here. For present purposes, my claim is only that any meaning-realist is at least committed to the existence of impure analyticities. For further discussion see Boghossian, forthcoming.

since, as Block observes, inferential role semantics are typically proposed as theories of narrow rather than of wide content. But is it true?

Not, I think, in any sense that matters to the issues presently at stake. It goes practically without saying that if we stipulate that what's in play is a notion of meaning that is not truth-conditional, then, relative to this stipulation, there won't be any sentences that are true purely by virtue of meaning, nor any inferences that are valid purely by virtue of meaning, truth and validity being—unavoidably—truth-theoretic properties. This would be, however, an entirely trivial consequence of the initial stipulation.⁴

It follows, therefore, that if we allow for a kind of meaning that is not truth-theoretic in nature, that a mere commitment to determinate meanings will not by itself entrain a commitment to analyticity—defined as truth (or validity) by virtue of meaning.

But that is irrelevant to the essential core of the argument I was concerned to develop and defend. The important point for my purposes—and the one that really engages Quine's critique of analyticity—is the observation that a believer in determinate meanings has no choice but to recognize that certain sentences or inferences will *hold*—to use as neutral a term as possible—merely as a result of the meanings of their constituent expressions. Whether this notion of 'holding' is cashed in terms of truth or assertibility or whatever—and, hence, whether the associated notion of analyticity is thought of as 'truth by virtue of meaning alone' or 'assertibility by virtue of meaning alone'—is immaterial from the standpoint of the argument. How it is cashed will depend, obviously, on which notion of meaning is in play and what conception of cognitive virtue it defines.

What Quine was concerned to deny—and my argument to affirm—was that there could be cases where the obtaining of the relevant cognitive virtue could be attributed solely to facts about the underlying meaning. Quine found it natural to think of that cognitive virtue as truth, and so framed the debate as an issue about the possibility of truth (or validity) by virtue of meaning alone. Clearly, however, had he been thinking in terms of assertibility, say, instead of truth, he would have been equally unhappy with the idea of 'assertibility purely by virtue of meaning alone'. What is essential to Quine is the claim that cognitive virtue, however defined, is never cogently seen as a consequence of meaning alone. And that is what I have argued a meaning realist can never accept.

F & L emulate Quine in thinking of meaning and, hence, of analyticity, truth-conditionally; in discussing their argument, I have followed suit.

⁴ And certainly not the sort of deep rejection of analyticity that Quine was after. Quine is very careful in 'Two Dogmas' to make clear that his rejection of analyticity is supposed to hold for a truth-conditional notion of meaning. (I don't mean to suggest that Block isn't aware of this; his point isn't meant to be a defense of Quine, but merely an attempt to stem the generality of my argument. I do not believe that it succeeds in that task.)

However, none of the crucial points depend upon this; they can all be recast, if necessary, in terms of a non-truth-conditional notion of narrow meaning. Since nothing hangs on it, I shall continue to use F & L's truth-conditional framework, leaving the generalization to other frameworks implicit.

6. *The Plausible-Sounding Principle*

Before we go back to the main line of argument, I have another bone to pick with Block. Block formulates a thesis he calls the 'Plausible-Sounding Principle' (this issue, p. 11):

Inferences that constitute an expression's meaning what it does are analytic.

He goes on to claim (a) that this principle is presupposed by F & L's argument, (b) that it is an instance of the more general connection claimed by me between meaning-realism and analyticity and, therefore, (c) that it is true for wide meanings but false for narrow meanings for the reasons that we had occasion to review and reject above.

I won't take a stand here on whether F & L really do require the PSP. But it is important to see that neither of Block's other two claims is true. The PSP is neither an instance of the general connection between meaning-realism and analyticity that I have outlined nor is it independently true, even for wide meanings.

It can seem as if the following argument simply clinches the PSP.⁵ If I say that the inference from 'rattling' to 'dangerous' is part of the MIRs of 'rattling' and 'dangerous', then I am saying that it is *necessary* for 'rattling' and 'dangerous' to mean what they do that that inference be valid. And if that inference's being valid is *necessary* for the expressions to mean what they do, then the expressions' meaning what they do is *sufficient* for the validity of the inference—which is just to say that the inference is analytic.

The fallacy here consists in conflating *being regarded as valid* with *being valid*. To say that the inference from 'rattling' to 'dangerous' is part of the MIRs of 'rattling' and 'dangerous' is only to say that someone has to be prepared to infer from 'rattling' to 'dangerous', if he is to mean by those words what they in fact mean. At best, this commits one to saying that someone has to be prepared to *regard* that inference as valid; it makes no commitment whatever on whether that inference is valid, let alone whether it is valid by virtue of meaning alone.

The point here is a general one. A theory of meaning may well end up

⁵ If truth be told, I myself argued in this way in an earlier draft of this paper.

deciding (and many theories of an 'Interpretational' sort have decided) that if someone doesn't believe certain obvious, pervasive *empirical* truths about dogs, then they can't mean *dog* by 'dog'. I am not saying that a theory of meaning is well-advised to decide this, just that it can do so without fear of logical inconsistency. To say that someone must regard a certain sentence as true, if he is to mean this by that, is not the same as saying that the sentence *is* true. *A fortiori*, it is not to say that the sentence is true by virtue of meaning alone.

Just so, an IRS theorist may end up ruling that no one counts as meaning *dog* by 'dog' unless that person is prepared to infer from 'x is a dog' to 'x is an animal'. But it doesn't *follow* that that inference is analytic. It may well be part of an inferential conception of meaning that all sorts of inferences that aren't valid by virtue of meaning alone are nevertheless parts of MIRs, nevertheless *constitutive* of meaning.

Notice that this point, too, is independent of the narrow/wide distinction. Even if we operated within a non-truth-conditional framework—and so thought of inferences as *holding*, *say*, rather than as *valid*—we would still need to recognize a distinction between *regarding* an inference as holding and its actually holding. Being constitutive of meaning is one thing, being analytic another.

In sum, then, I reject both the claim that the PSP is a special case of the internal connection between meaning-realism and analyticity and, indeed, the PSP itself.

7. *F & L Continued*

Returning to the main line of argument, let us take stock. F & L have argued that an IRS can't identify the MIR of an expression with its role in analytic inference, because the analytic/synthetic distinction is 'unprincipled'. According to their official explanations, for an inference to be analytic is for it to have a purely semantical property: validity by virtue of meaning alone. And for the *a/s* distinction to be unprincipled is for there to be no inferences with such a semantical property, no inferences that are analytic in the purely semantical sense.

As against this, I have claimed that no meaning realist — *a fortiori*, not F & L—can agree that there will be no analytic inferences in the semantical sense. So their official argument against the possibility of a non-holistic IRS fails.

At this point it might have been appropriate for me to elaborate on the entailment I have claimed; and I would gladly have done so were it not for the fact that—curiously enough—F & L give every indication of accepting it! Thus, they say (p. 337, fn. 3):

There are causal theories of meaning around And like any account of meaning, each of them implies a corresponding notion of analyticity.

Their subscription to the entailment is further fleshed out in the book from which the article currently under discussion is drawn:

[I]magine somebody who holds a reductionist theory of meaning; for example, somebody like Skinner, who holds that for 'dog' to mean *dog* in a certain speaker's mouth is for the speaker to have the habit of uttering 'dog' when there are dogs around. . . . Look, someone might say, if Skinner were right, then you could have a situation in which a speaker has two responses ('dog' and 'shmog', as it might be) that are conditioned to *exactly the same stimuli*. But then it would follow that these responses would be *synonymous* for that speaker. So then the following sentence would be *analytic* in the speaker's language. . . . 'Whatever is a dog is a shmog.' But Quine showed that there are no such things as synonyms or analytic sentences. So Skinner's semantics must be wrong. *A priori!* In fact, *all* semantic theories must be wrong, *a priori*, except for the nihilistic theory which says that there are no semantic properties.

(Fodor and LePore, 1992, pp. 56–7)

F & L rightly point out—just as I do in 'Analyticity' (Boghossian, forthcoming)—that this pattern of reasoning is a *reductio* of the claim that there can be no analytic inferences in the purely semantical sense.

Their official exposition notwithstanding, then, their argument, against the targeted IRS theorist can't after all be that it is committed to the existence of inferences that are valid purely by virtue of meaning; by their own admission, every theory of meaning is, including the Skinner-style informational theory that they are concerned to promote. So *what* is their argument?

It's hard to find a different argument for the same conclusion in the article that is the subject of this symposium. But in their book, F & L provide the following alternative explanation of what Quine may be credited with having shown (1992, p. 57):

We take it that, strictly speaking, Quine in 'Two Dogmas' did *not* show, or even argue, that there are no analytic truths.

'He did argue', they continue in a footnote,

plausibly in our view, that you can't reconstruct analyticity by appeal to apriority; that is, that if there *are* analytic truths, we don't know them *a priori*.⁶

⁶ But compare p. 25 of the same book (Fodor and LePore, 1992):

When Quine says, 'No *a/s*' he presumably means 'No analytic sentences'.

It is unfortunate that the argument of this interesting book is marred by this sort of vacillation about what 'No *a/s*' is supposed to mean.

So maybe what F & L really intended to say is that a non-holistic IRS is committed to *analyticities that are a priori knowable* and that *that* is the notion that has been shown to be empty by Quine.

Well, if that is the argument F & L wished to make in the paper that is the subject of this symposium, it is arguable that they should have expressed themselves quite differently; but let us put that to one side. It is a substantially different argument from the one we have just considered and is not subject to the same reply. The point is that a priori analyticity is a much more demanding notion than a plain semantical analyticity; it cannot be defended by mere appeal to meaning-realism—it simply doesn't follow from a commitment to determinate meanings that meaning equivalences will be knowable in any particular way, let alone a priori. How, if at all, facts about the relations between meanings will be knowable will depend on the detailed nature of meaning, detail on which a mere realism about meanings is silent, and which awaits resolution by a satisfactory theory of meaning.

For all that, could it reasonably be claimed, as F & L do, that Quine showed that there couldn't be any a priori analyticities, so that no satisfactory theory of meaning could afford to postulate them? In my opinion, not at all (here again see my discussion in 'Analyticity' (Boghossian, forthcoming)).

However, I don't need to argue the point here because it is completely irrelevant to the issue at hand. The fact is that *it is simply not true that an inferential role semantics is committed, by its very nature, to reconstructing meanings out of inferences whose validity is knowable a priori*. So even if we were to grant that Quine had shown that the extension of 'a priori analytic inferences' were empty—something which for independent reasons I don't believe we should be prepared to do—F & L's argument against an IRS fails.⁷

To get an argument that works, relative to the concession to Quine, F & L would somehow need to show that any IRS is committed to the inferences' being a priori. But where would such a demonstration come from? The only basis I can think of would have to consist in some general argument to the effect that *any* meaning-based inference has to be a priori. Whatever the merits of such a claim, however, it is not one that F & L can appeal to, for the theory of meaning that they favor is knee-deep in analytic inferences that aren't a priori: as they themselves are fond of pointing out, causal/informational theories of meaning will generate lots of inferences whose validity is secured by meaning alone, but which will not be knowable a priori. So they can hardly believe that it is a general

⁷ I do not deny that there are actual IRS theorists who thought that the MIRs defined by their theories would contain only a priori knowable inferences; some of them may have even thought of that as a virtue. But F & L purport to have an argument against *any* sort of inferential role semantics. And that is the argument that is eluding my grasp.

requirement on analytic inferences that they be a priori. If there were such a requirement, it would count against a very wide range of theories of meaning, including their own (always relative, of course, to the concession to Quine).

I conclude, therefore, that regardless of whether 'analyticity' is read as expressing pure semantical analyticity or as expressing a priori analyticity, F & L's argument discloses no new difficulty for an inferential role semantics. It remains possible, of course, that there is some *other* notion of analyticity which is such that it is both required by an IRS and effectively destroyed by Quine. But if there is such a notion, I don't know what it is. And so far as I am able to judge, neither do F & L.

Addendum

In their reply to this piece, Fodor and LePore introduce a *third* thesis that they attribute to Quine and which they label a 'rejection of the analytic/synthetic distinction'. This thesis is neither the one that they defined in the paper that my own piece responds to—*viz.* that there are no sentences that are true purely by virtue of what they mean. Nor is it the thesis that they attribute to Quine in their book on *Holism*—*viz.* that there aren't any *a priori* knowable analyticities. Rather, it is the thesis that there are (this issue, p. 43):

no analytic inferences in the operative sense (viz. . . . no inferences whose acceptance is constitutive of your meaning what you do by what you say.

There are many things I am tempted to say about this new suggestion, but in this short note the following will have to do. If that really is the thesis that F & L think Quine has demonstrated, then their argumentative strategy is even more puzzling than my previous description of it allowed. For suppose you believe the thesis that they say they believe. And suppose you wish to argue against an inferential role semantics. Then, piece of cake (as we say in the States).

1. An IRS just *is* the view that what you mean by your words is constituted by (some) of the inferences that you accept. (Definition)
2. There are no inferences whose acceptance is constitutive of what you mean by your words. (Thesis allegedly shown by Quine)
3. Hence, an IRS is false.

It is really as simple as that. There is no need for the elaborate detour through compositionality. Nor for anything else. An IRS is committed by definition to there being meaning-constituting inferences; if Quine is *directly* taken to have shown that there aren't such inferences, or even if he is taken to have shown simply that there is no fact of the matter as to which inferences are meaning-constituting and which ones aren't, then an IRS is sunk. Period. Why, if that is what they meant, didn't F & L simply say *that*?

The answer, I suspect, is that they didn't have this immediately question-begging interpretation of 'no analytic/synthetic distinction' in mind when they wrote their original piece; else they wouldn't have written it. What made F & L's argument seem interesting is that it promised to show that an inferential role semantics was committed, somewhat unobviously, to a notion—truth by virtue of meaning—that some people had been made to feel uncomfortable about as a result of reading Quine. As I argued above, that argument fails; but at least it's an *interesting* argument. F & L's new argument is much worse than their old one: it simply *begs* the question against the inferential role semanticist.

*Department of Philosophy
New York University
New York, NY 10003
USA*

References

- Boghossian, P. forthcoming: Analyticity. In C. Wright and R. Hale (eds.), *Blackwell Companion to the Philosophy of Language*. Oxford: Blackwell Publishers.
- Fodor, J. and LePore, E. 1991: Why Meaning (Probably) Isn't Conceptual Role. *Mind and Language*, 6, 328–43.
- Fodor, J. and LePore, E. 1992: *Holism: A Shopper's Guide*. Oxford: Blackwell Publishers.