

INFERENCEAL ROLE SEMANTICS AND THE
ANALYTIC/SYNTHETIC DISTINCTION

(Received 15 September 1993)

1. INTRODUCTION

Towards the end of their provocative discussion of conceptual role semantics, Jerry Fodor and Ernie Lepore write:

If, as we suspect, Quine is right about the a/s distinction, then the moral of our discussion is that [a conceptual role semantics] is false. (p. 186)

For ease of reference, call this conditional claim of theirs *Falsity*. As we shall see, *Falsity*, and the argument for it that F&L deploy, embed many of the issues that are most central to the book that is the subject of this symposium. Furthermore, since, apparently, many philosophers are inclined to believe that Quine was right about the a/s distinction, and since, apparently, an equally large number of them are inclined to believe in some form of conceptual role semantics, F&L's claim is of great independent interest. I will argue that *Falsity* is false. And I will attempt to draw some morals from its falsity. But first some preliminaries.

To begin with, although the cited paragraph talks about "conceptual role semantics", Fodor and Lepore's argument concentrates exclusively on an expression's *inferential* role – i.e., on the purely linguistic aspect of an expression's overall conceptual role, a role that will typically include relations to non-linguistic stimuli or objects as well. F&L's idea is that, if Quine is right about the a/s distinction, then the purely inferential part of any conceptual role semantics must be deemed hopeless, and hence so must any conceptual role semantics. Henceforth, therefore, I will talk exclusively about F&L's argument against an inferential role semantics (hereafter, IRS).

Philosophical Studies 73: 109–122, 1994.

©1994 Kluwer Academic Publishers. Printed in the Netherlands.

What, then, is an inferential role semantics, precisely? Clearly, from any particular sentence, I will be prepared to infer to some other sentences, and not to others. For example, from

- (1) x is a dentist.

I am prepared to infer that

- (2) x makes a living by tending to other people's teeth.

I am also disposed to infer that

- (3) x is probably rich.

By contrast, I am not prepared to infer that

- (4) x snores loudly at night.

Let's call the totality of the inferences to which a sentence is capable of contributing its *inferential role*. A word'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.

Now, one question that has obsessed philosophers in recent times has been the following: By virtue of what sort of fact about the word "dentist" does it mean whatever it in fact means, say, *dentist*? An inferential role semantics is an answer to this sort of question. Its idea is that meaning facts are facts about inferential role: an expression means what it does by virtue of participating in one set of *inferences* rather than another. We may put this by saying that, according to an IRS, there is some construct out of an expression's total inferential role that constitutes its meaning what it does. Call this construct an expression's *meaning-constituting inferential role*.

If expressions mean what they do by virtue of the inferences they participate in, then some inferences are *constitutive* of an expression's meaning what it does, and others aren't. And the pressing question is: Which are which? What property does an inference have to have, if it is to be meaning-constitutive?

All the participants to the present debate agree that to this day no one has succeeded in providing a systematic answer to these sorts of questions. As yet, there are no plausible accounts out there of what properties an inference must have if it is to be meaning-constitutive.

2. F&L ON ANALYTICITY

F&L's interesting claim, however, is that no such account is forthcoming, if Quine is right about there being no analytic/synthetic distinction. How is this argument supposed to work?¹ The rough idea is this. First, F&L assume the following principle – call it *The Principle*:

If someone's being disposed to make the inference from S1 to S2 is *constitutive* of what S2 means in that person's idiolect, then the inference from S1 to S2 is *analytic* in that person's idiolect – i.e., valid by virtue of the meanings of S1 and S2 alone.

Using *The Principle*, they argue that an IRS is committed to the existence of analyticities. This part is quite straightforward: by definition, an IRS is committed to the existence of meaning-constitutive inferences, and by *The Principle*, those inferences are analytic. F&L's strategy is then to show that "Quine's being right about the a/s distinction" is somehow in conflict with the truth of that commitment. What we need to do is find out first, what precisely it means to be committed to the existence of analytic inferences, and second, how it is that Quine's being right about the a/s distinction conflicts with its truth.

Beginning with the first question, what is it for an inference to be analytic? Here F&L's definition is reassuringly familiar: for an inference (or sentence) to be analytic is for it to be valid (or true) by virtue of its meaning alone:

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. (p. 179)

But what exactly does it mean to say that a sentence is *true by virtue of meaning alone*? What sort of semantical fact does this description purport to describe? Here there lurks, I believe, an ambiguity, one that it will be important for our purposes to expose.

The intuitive idea does not seem difficult to state: a sentence is true by virtue of its meaning alone when it is such that: facts about its meaning suffice for its truth all by themselves, without any contribution from any other fact. What calls for further clarification, however, is the final phrase: without any contribution from any other fact.

Clearly, the phrase is intended to exclude a dependence on ordinary contingent worldly facts, such as, for example, that cats scratch furniture or that snow is white. But is it really intended to be taken completely literally, excluding even a dependence on such *logical facts* as that everything is identical with itself? Of course, if the truths of logic were themselves mere truths by virtue of meaning alone, there would be no genuine contrast here. But suppose they aren't. Then does a sentence not count as analytic if its truth depends irreducibly on a truth of logic?

Obviously, there are two quite distinct notions here, between which the literature on analyticity has not always properly distinguished. (Indeed it's arguable, though I won't argue it here, that Quine himself does not always keep them neatly apart.) Let us say that a sentence is *purely analytic* just in case it is such that its truth depends on its meaning alone, without any irreducible contribution from anything else, including logic. And let us say that it is *impurely analytic* if its truth depends on its meaning plus logic, but without any contribution from any further sort of fact.²

The most important advocates in recent times of the idea of pure analyticity were, of course, the positivists. It is presupposed by one of the most important doctrines of their philosophy – namely, by their claim that facts about meaning alone explain where the truth of the sentences of logic and mathematics come from. Clearly, if the idea of truth by virtue of meaning alone is to *explain* where logical truth comes from, then it cannot presuppose it. It follows, therefore, that any claim to the effect that logical truth is analytic truth must be understood as framed in terms of a notion of *pure* analyticity – it must be the claim that logical truth is *pure* analytic truth. The claim that logical truth is merely impurely analytic is of no interest; it is true trivially that the truths of logic are true by virtue of their meaning plus logic.

Pure analyticity, then, is a *very* demanding notion: for something to

be purely analytic it must be true by virtue of its meaning alone and without any irreducible contribution from anything else, not even from such a simple truth of logic as the law of self-identity. By contrast, the notion of impure analyticity is a much less ambitious idea. Since impure analyticity *presupposes* an unreconstructed notion of logical truth, it cannot purport to provide a philosophy of logic. What it can do, much more modestly, is mark out a class of sentences for special consideration, on the grounds that they are transformable into the logical truths by, for instance, the substitution of synonyms for synonyms. How interesting or important such a classification will turn out to be will depend on how interesting or important the notion of logical truth is. From the standpoint of the impure notion, however, that is a separate question, one that it is not in the business of addressing. Quine, of course, was *equally* opposed to the idea of impure analyticity, and in his most famous paper on the subject, "Two Dogmas of Empiricism," he launched an attack on it as well.

Now, F&L do not make a distinction between pure and impure analyticity; but I take it that in their arguments they must have the impure version in mind (for example, their references are uniformly to "Two Dogmas" which is practically exclusively concerned with the impure notion). At any rate, it's certainly the notion they *should* have had in mind: there is absolutely no reason to believe that an inferential role semantics is committed, just as such, to as heady a notion as that of pure analyticity.³ Eventually I will be arguing that The Principle is false even when it is read as claiming that meaning-constituting inferences are *eo ipso* impurely analytic. However, it certainly has no chance of being true if it is read asserting that a meaning-constituting inference is *eo ipso* purely analytic, as if a semantic theory could by itself ground something as potent as the linguistic doctrine of logical truth. For most of the rest of this paper, then, I will assume that the notion of analyticity that's at issue is that of impure analyticity.

We next need to face the question: What, according to F&L, is it for Quine to be right about there being no *a/s* distinction (with "analyticity" understood as standing for the impure notion)? Here, unfortunately, matters are less than clear-cut. In particular, F&L attribute at least two distinct theses to Quine, between which we must distinguish if we are

to assess their argument. The first thesis occurs in the chapter presently under discussion:

There are, as his critics have often remarked, two ways of reading Quine's claim that there is no *a/s* distinction: either that the distinction is *incoherent* or that the distinction is coherent but that *the class of analyticities is empty*. We've argued that, given the compositionality of meaning, neither of these readings of Quine's rejection of the *a/s* distinction is compatible with [IRS]. (p. 183)

This seems fairly clear: to say that Quine was right about the *a/s* distinction is to commit oneself either to the thesis that the distinction is incoherent (whatever precisely that comes to) or to the thesis that there are no analyticities.

However, in an earlier chapter on Quine we find the following passage:

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. (p. 57)

Now this seems to me to attribute an entirely different accomplishment to Quine. Here Quine is said not only not to have shown that the class of analyticities is empty, but not to have even tried to show that. Rather, the claim is that he merely refuted a certain *theory* of analyticity, one that attempts to reconstruct it in terms of apriority (whatever precisely that comes to).

Here, then, is how I propose to proceed. Putting aside exegetical concerns, I will distinguish between the two principal theses that are attributed to Quine in the above passages: that the class of analyticities is empty – call this *Emptiness* – and that analyticity can't be reconstructed in terms of apriority – call this *Non-reconstruction*. I will then argue that their argument for Falsity doesn't work with either thesis.

3. THE ARGUMENT FOR FALSITY – EMPTINESS

Falsity, you may recall, is the view that if Quine is right about the *a/s* distinction, then an inferential role semantics is false. If we read Quine's

being right about *a/s* as consisting in the emptiness of the class of impure analyticities, then with The Principle in place, Falsity is secured very quickly. For if Quine is right, then there are no analyticities. But, according to an inferential role semantics, an expression has a meaning only to the extent that it participates in some meaning-constituting inferences. And according to The Principle, any meaning-constituting inference is *eo ipso* analytic. Hence, if there are no analyticities, then an inferential role semantics is false.

Now, I think that there is no doubt that this argument is valid. The trouble with it is that it proves too much. Given Emptiness – the claim that there are no analyticities whatever – and even without using The Principle, we can prove not merely that an inferential role semantics is false, but that *all theories of meaning are false*.

To see why, we need to look briefly at a train of thought that I have developed at greater length elsewhere (again, see “Analyticity”). 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, as we must, there are bound to be facts about whether two expressions do or do not mean the same – there are bound to be facts about synonymy. But an inference will be impurely analytic just in case it is transformable into a logically valid inference by the substitution of synonyms for synonyms. It follows, therefore, that if meaning realism is true, then there are bound to be facts about whether the inference from

x is a vixen

to

x is a female fox

is or is not impurely analytic.

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 my forthcoming.)

Making a long story very short, then, I don't see how there could fail to be inferences that are at least impurely analytic, if meaning realism is true. Hence, if Emptiness is true, and there are no analyticities anywhere, meaning realism must be false.

If, then, F&L's argument for Falsity *does* appeal to the alleged emptiness of the class of impure analyticities (which is actually what they *say*) then their argument doesn't work. For F&L's purpose isn't to argue that there are no meaning facts whatsoever, but rather that an inferential role semantics provides the wrong account of such meaning facts as there are. So far, however, we have failed to uncover an argument that would have only that thesis as its conclusion.

4. THE ARGUMENT FOR FALSITY – NON-RECONSTRUCTION

I will next turn to asking whether their argument for Falsity fares any better with the alternative construal of Quine's rejection of an *a/s* distinction that they also mention in their book. As we saw, this construal is formulated in terms of Non-reconstruction: It's not that there aren't any inferences that are valid by virtue of meaning alone; it's simply that their analyticity can't be reconstructed in terms of their apriority.

Well, what precisely does this mean? As best I can make out, the idea seems to be that Quine showed that there is no necessary equivalence between the property of being analytic and the property of being an inference that's held on a priori grounds, independently of experience (see pp. 57–58).

Ignoring many complications in how all this is to be fleshed out, we may present the following sketch of an argument for Falsity. According to an inferential role semantics, some inferences are constitutive of the meanings of the participating sentences – that is, someone's being prepared to make those inferences is necessary for those sentences to mean what they do in his idiolect. According to *The Principle*, these meaning-constituting inferences are analytic. And according to Non-reconstruction, it is not possible to specify the analytic inferences by

saying that they are the ones that are held come what may, i.e., a priori. Therefore, if Quine is right about the a/s, then an inferential role semantics is false.

Now, whereas the argument using the Emptiness of impure analyticities was valid but too strong, the present argument is clearly invalid. It is missing a premise to the effect that, if the analytic inferences can't be specified by reference to the inferences that are a priori, then there will be no other way of specifying them available to an inferential role theorist. But I know of no argument for this missing premise; nor do I find one in F&L. Why can't the IRS theorist say: "Oh, I see. Quine showed that you can't reconstruct analyticity in terms of apriority. Given The Principle, I suppose that shows that I can't specify the inferences that are meaning-constituting by reference to the ones that are a priori. Good, I'll have to think of something else." What F&L need is some argument to the effect that the *only* way available to an IRS theorist for specifying the extension of 'analytic inference' is by recourse to apriority. But I don't see that they have such an argument.

As I conceded at the outset, to this day no IRS theorist has proposed a satisfactory way of specifying the inferences that are supposed to be meaning-constituting. Everybody is aware of this fact. What we were promised by F&L, however, was a principled argument to the effect that if some allegedly widely believed thesis about the a/s distinction were true, then there would be something like a proof that no inferential role semantics could work. Unfortunately, however, we have failed to uncover any such proof.

I have looked at two arguments for Falsity, one appealing to Emptiness and the other to Non-reconstruction, and without even challenging their premises I have found them wanting. However, in my view, the most interesting place at which F&L's argument goes wrong – both in their discussion of Falsity and, as we shall see, in the central argument that drives their book as a whole – rests in their appeal to The Principle. I think that The Principle, plausible as it has seemed to many ears, is actually false. And quite apart from the present context, it seems to me both quite important that The Principle is false and insufficiently appreciated that it is. In the remainder of this paper I would like to start looking at why.⁴

5. THE PRINCIPLE

According to The Principle, the constitutivity of an inference entails its analyticity:

If someone's being disposed to make the inference from S1 to S2 is *constitutive* of what S1 means in that person's idiolect, then the inference from S1 to S2 is *analytic* in that idiolect – i.e., valid by virtue of the meanings of S1 and S2 alone.

Now, there is no doubt that The Principle sounds plausible. Certainly, many writers in the philosophy of language have assumed it, often without any explicit argument. Usually, it has turned up as a principle about *sentences* rather than inferences, as follows:

If someone's being disposed to hold S true is *constitutive* of what S means in that person's idiolect, then S is analytic in that idiolect – i.e., true by virtue of its meaning alone.

The extent to which philosophers have been inclined to believe some version or other of The Principle has manifested itself in the puzzlement that many have felt at Quine's view on the philosophy of logic, at least as these are developed in "Carnap and Logical Truth" and in *Word and Object*. Quine argued both that it was wrong, or worse, to say that logic was true by virtue of the meanings of the logical constants (remember, this would have to be understood as the claim that they are purely analytic); and also that we can't make genuine sense of the possibility of an *alternative logic*, so that if we encountered someone using a set of logical principles that were contrary to our own, we would most likely have to say that

he is merely using the familiar particles 'and', 'all', or whatever, in other than the familiar senses, and hence that no real contrariety was present after all.⁵

Quine makes similar remarks about the suggestion that there could be a 'pre-logical' people – people who are disposed to accept certain simple self-contradictions. Quine says very strongly:

We are left with the meaninglessness of the doctrine of there being pre-logical peoples; pre-logicality is a trait injected by bad translators. This is one more illustration of the inseparability of the truths of logic from the meanings of the logical vocabulary. (p. 102)

What Quine is insisting on in these passages is the constitutiveness of certain sentences of logic: If a person is to mean conjunction by 'and', he must hold certain sentences involving 'and' true. However, he maintains, it does not follow that the truths of logic are true by virtue of their meanings alone. This had led many philosophers to wonder how Quine could possibly expect to have it both ways. How could it both be the case that the logical truths aren't true by virtue of their meaning alone, and also such that, giving them up would immediately result in a change in their meaning? Aren't these two claims simply incompatible with one another? (Notice that this is *not* a question either about the plausibility of the Quinean claim about alternative logics, nor a question about the plausibility of Quine's rejection of the analyticity of logic; rather, it is exclusively a question about the *compatibility* of these two claims.)

The claims *would* be incompatible, of course, if The Principle were true. An I think the fact that it seems to be so easy for people to hear them as incompatible is a reflection of the grip that The Principle appears to have over philosophers. But is The Principle true?

Well, why doesn't the following simple argument simply clinch The Principle? If I say that holding true the sentence 'All dogs are animals' is constitutive of 'dog' 's meaning *dog*, then I am saying that it is necessary for 'dog' to mean *dog* that that sentence be held true. And If I say that that sentence's being held true is necessary for the expression to mean what it does, then I am saying that the expression's meaning what it does is *sufficient* for the sentence's being held true – i.e., that it is analytic.

The fallacy behind this argument for The Principle occurs in the very last step. The point is that there is all the difference in the world between saying that a certain sentence must be *held true*, if it is to mean this, that or the other, and saying that it *is* true. *A fortiori*, there is all the difference in the world between saying that it must be held true, if it is to mean this, that, or the other, and saying that it is *true by virtue of its meaning alone*. To claim that a certain sentence is constitutive is just to say that it must be held true; it is not to say anything, in and of itself, about whether the sentence is true. *A fortiori*, it is not to say anything about its being true by virtue of its meaning.

Similarly for inferences. To say that the inference from 'red' to 'not blue' is constitutive of the meaning of 'red', is to say that someone must be prepared to make that inference if he is to mean *red* by 'red'. If we wish, we may put this by saying that such a person must be prepared to *regard* the inference from 'red' to 'not blue' as valid. But this doesn't *by itself* imply that the inference is valid, let alone that it is valid by meaning alone. Being constitutive of meaning is one thing, being analytic another. The Principle is false.

Still, one may legitimately wonder whether there isn't a reason short of entailment that could motivate the claim that all constitutive inferences are analytic. After all, what *other* basis could a theory of meaning have for ruling a certain sentence or inference constitutive of meaning, if not its presumed analyticity?

As Quine noticed, there could be other such bases. In the case of the logical principles, their constitutive status is "adequately accounted for," he wrote,

by the mere obviousness of logical principles, without help of a linguistic doctrine of logical truth. For, there can be no stronger evidence of a change in usage than the repudiation of what had been obvious, and no stronger evidence of bad translation than that it translates affirmations into obvious falsehoods. (p. 106)

Constitutiveness might be motivated by mere obviousness: some empirical truths might be so obvious that, someone's giving them up would always be better explained by the hypothesis that their meaning has changed than by the hypothesis that they are now, mysteriously enough, believed false. Whether constitutiveness might be motivated in some further way is a question that I have to postpone for another occasion. For our purposes the preceding suffices. We are already in a position to see that constitutiveness neither entails analyticity, nor is exclusively explainable by it. If this is right, then various interesting consequences follow.

First, and most obviously, we see that the question whether an inferential role semantics is viable is wholly independent of the issue about analyticity. In particular, we see that someone may safely propose an inferential role account of the meanings of the logical constants without committing himself to a heady and possibly confused doctrine of the pure analyticity of logical truth. Indeed, we see that an inferential role

semantics might be right even if nothing is even impurely analytic. Of course, as I have already argued, I believe that any meaning realist is committed to the existence of at least some impure analyticities, so my motivation for emphasizing this fact doesn't stem from my pessimism about impure analyticity. However, even if I were wrong about this, it still wouldn't follow that an inferential role semantics was false.

A fortiori, and this is my second point, it wouldn't follow from the fact that there are *too few* impure analyticities that an inferential role semantics was false. I suppose that many philosophers have thought that Quine went overboard in occasionally asserting that there are no impure analyticities whatever, that the correct view is that there are many fewer (impure) analyticities than had traditionally been thought. Putnam, I believe, has been the most prominent advocate of such a view. Nothing in my argument connecting meaning realism with analyticity speaks to the question how many analyticities there are, so for all that that argument shows, Putnam may well be right. Now, if The Principle were true, and there were only a very few analyticities, then this would itself pose a significant threat to an IRS. For, given The Principle, an IRS is committed to distinguishing between meanings only by exploiting analyticities, and if there are no analyticities that distinguish between, say, 'gold' and 'silver', or 'cat' and 'dog', then an IRS is bound to be false. However, as I have argued, The Principle is false and an IRS is not restricted to appealing to analyticities in its account of what inferences count as meaning-constitutive.

My third and final point is central to the basic concern that motivates much of F&L's book. Much of their work is driven by the fear that, if at least *some* of an expression's inferential liaisons are allowed to be constitutive of its meaning, then a slippery slope argument will lead to the conclusion that they all are, and horrible holism will result. This is the ur-argument for meaning-holism that they outline in Chapter 1, and which they call 'Argument A'. Such is their confidence in the slipperiness of this particular slope, that they spend their entire book trying to undermine its very plausible-seeming first premise.

But they needn't have been so concerned. The point is that the only reason they provide for thinking that there is such a slippery slope, depends on the non-existence of a distinction between the analytic and

the synthetic. However, if The Principle is false, then the question whether there is a distinction between the analytic and the synthetic is simply orthogonal to the question which inferences are constitutive of meaning, and so what reason there may have been to think that there is a slippery slope vanishes.

Of course, there still remains the question – fundamental for an IRS – how the difference between the meaning-constituting and the non-meaning-constituting inferences is to be specified. And as I said at the very outset, that remains an unsolved problem. Indeed, for all I know, it can't be solved, because for all I know meaning does not supervene on inferential role. But if it can't be solved, it won't be because of some problem about analyticity that has been newly brought to our attention by F&L.⁶

NOTES

¹ In a paper that is derived from their book, entitled "Why Meaning (Probably) Isn't Conceptual Role," *Mind and Language*, Fall 1992, Fodor and Lepore present a somewhat different argument. I discuss that argument in my "Does an Inferential Role Semantics Rest Upon a Mistake?" in *Mind and Language*, Spring 1993.

² For further discussion of all this see my "Analyticity" forthcoming in Hale and Wright (eds.): *A Companion to the Philosophy of Language* (Oxford: Blackwell, 1994).

³ For reasons of space, I have to be crude here. My own view is that, although Quine was right to be suspicious of pure analyticity, his reasons for the suspicion are not correct. I explain this in "Analyticity."

⁴ Here, again, the issues are a lot more complicated. For a fuller treatment see "Analyticity".

⁵ "Carnap and Logical Truth," reprinted in *The Ways of Paradox* (New York: Random House, 1966), p. 102.

⁶ I am grateful to Barry Loewer and Ned Block for comments on an earlier draft of this paper and to Stephen Schiffer for helpful discussion.

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