

Article

Holism, Hyper-analyticity and Hyper-compositionality

NED BLOCK

1. *The Old Hat*

Fodor and LePore claim to have a new argument against inferential role semantics, a view that they regard as gospel not only in linguistics and philosophy, but also throughout cognitive science. I regard inferential role semantics as an embattled minority position that deserves to be much more widely accepted than it is.

Before I try to defend inferential role semantics against Fodor's and LePore's new argument, I want to make some points about the old one, which I shall call the Old Hat.¹

To begin with, inferential role semantics is the doctrine that meaning is inferential role. Supposing that one knows what it means to speak of

This paper is a reply to Fodor's and LePore's (1991), 'Why Meaning (Probably) Isn't Conceptual Role', which is a version of their critique of Block (1986), Chapter 5 of their (1992): 'Ned Block: Holism and Conceptual Role Semantics'. Though I have both versions in mind, I am keying my responses mainly to the paper in this journal (1991). Page references should be understood as indexed to that work unless (1992) is specifically cited.

An earlier version of this paper was given in response to Fodor and LePore at the meeting of the Society for Philosophy and Psychology, June 1991, and the July, 1991 meeting of SOFIA in Salamanca. I am grateful to Paul Boghossian, Martin Davies, Paul Horwich, Jerry Katz, Eric Lormand and Robert Stalnaker for comments on an earlier draft, and also to Jerry and Ernie for our lively discussions. I am also grateful to Earl Conee for sending me his reply to Fodor and LePore, which takes a very different line from mine. Address for correspondence: Department of Linguistics and Philosophy, M.I.T. (20D-213), Cambridge, MA 02138, USA. Email: block@psyche.mit.edu.

¹ See Fodor, 1987, pp. 55-6; Block, 1986, p. 629.

2 *Mind & Language*

the inferential role of a sentence,² what does it mean to talk of the inferential role of a word? In the first instance, it is sentences, not words, that one infers from and to, and thus it is sentences that *have* inferential roles. Nonetheless, we can think of the inferential role of a word as represented by the set of inferential roles of sentences in which it appears. The notion of the inferential role of a word (and likewise for other sub-sentential constituents) is a made-up idea, but this is an obvious way to make it up.

Here is the Old Hat. Suppose that I infer from 'This is a rattling snake' to 'This is dangerous',—for short, I will speak of inferring from 'rattling snake' to 'dangerous'. Now if we include *all* of a sentence's inferential liaisons—that is, all of the inferences in which a sentence participates—in its inferential role, then, carrying this decision through to words, the inference from 'rattling snake' to 'dangerous' will be part of the inferential roles of—and hence part of the meanings of—the words 'rattling', 'snake' and 'dangerous' (I will ignore 'this', 'a' and 'is'). If I infer from 'rattling snake' to 'dangerous', and you do not, then we do not share the same meanings of any of 'rattling', or 'snake', or 'dangerous'. If we don't share the meanings of some of these words, then the meaning of 'Rattling snakes are dangerous' that I accept isn't the same as the meaning of this sentence that you do not accept. But then how can people ever disagree? The same reasoning leads to the conclusion that I can't change my own mind, since if I accept a sentence, and then later reject it, the meaning of the sentence that I accept isn't the same as the meaning of the sentence that I later reject. Further, apart from the example of perfect Twins, people are certain to differ in their 'snake' beliefs or in beliefs that are inferentially related to those beliefs, or in Thus, the meaning or content that you assert with 'Rattling snakes are dangerous' is most unlikely to be the same as the meaning or content that I assert with the same sentence, so agreement is almost as problematic for the inferential role semanticist as disagreement.

This is the problem that leads so many to abhor holism, a doctrine whose incarnation in inferential role semantics is the claim—assumed above—that all the inferences in which an expression participates are included in its inferential role.

2. *The Crack*

As wonderful as the Old Hat is, it is not what Fodor and LePore come to tell us about. They regard the Old Hat as 'external to semantics', having to do rather with epistemology and ontology. They are after a new problem, a *linguistic* problem, the alleged Crack in the foundation of semantics. The

² This is something of a fiction, a useful one I think, on which both Fodor's and LePore's critique and my response, are based.

Crack involves an inconsistent triad: inferential role semantics, compositionality and the rejection of the analytic/synthetic distinction. Compositionality is non-negotiable, they say, and since no one wants to accept the analytic/synthetic distinction (or so they say), we inferential role semanticists are in deep trouble.

The argument goes like this: with no analytic/synthetic distinction to appeal to, how can the inferential role semanticist find a principled basis for distinguishing among inferential liaisons, including some in his inferential roles, while excluding others? He cannot postulate zero inferential liaisons for very many words, since that would give all those words the same 'zero' meaning. So, with no way to justify stopping between none and all, he must include all inferential liaisons in inferential roles. But then inferential roles are not compositional. For if we infer from 'rattling snake' to 'dangerous', then this inference is part of the inferential role of 'rattling snake'. This inferential role, being part of the meaning of 'rattling snake', is required to be a function of the inferential roles of 'rattling' and 'snake', by the principle of compositionality. But, according to Fodor and LePore, the meanings of 'rattling' and 'snake' *in fact don't determine* that rattling snakes are dangerous. So if all inferential liaisons are included in inferential roles, compositionality cannot be satisfied.

The only way out for the inferential role semanticist, say Fodor and LePore, is to give up holism and accept the analytic/synthetic distinction, for then we can nominate the analytic inferences as the ones that are *in* inferential roles, the synthetic ones being out. Analytic inferences, they say, are compositional. The upshot: the inferential role semanticist is caught between giving up compositionality on the one hand, and having the analytic/synthetic distinction forced upon him on the other.³

3. *The Crack Depends on The Old Hat*

I will offer a clarification, and then an objection. First, the clarification: Fodor and LePore seem to assume, along with many others, that Quinean objections to analyticity aside, the inferential role theorist has the option of appealing to analyticity as a way of discriminating the inferential liaisons that are in inferential roles from those that are out. But if we stick to traditional ideas about the extension of 'analytic', there aren't *enough* analyticities.

Consider the putative analytic truths involving 'cat'—'Cats are animals', 'Cats are living beings', 'Cats are grown up kittens', *etc.* The problem is that abstracting from the words 'cat', 'kitten', *etc.*, appearing in these sentences, there is nothing here to distinguish 'cat' from 'dog'. Corresponding to 'Cats are grown up kittens', we have 'Dogs are grown up puppies'.

³ The argument is neatly summed up in the last paragraph of their Section 2.

4 *Mind & Language*

Sure, 'Nothing is both a cat and a dog' can be used, but so can 'Nothing is both a dog and a cat'. Even if 'Cats are feline', and 'Dogs are canine' are analytic, this is of no help without other analytic truths that distinguish 'feline' and 'canine'. Jerry Katz (1972) puts forward a view according to which words like 'feline' and 'canine' could be used to distinguish between the meanings of 'cat' and 'dog', but 'feline' and 'canine' would have to be primitives. That may be fine for his purposes, but an inferential role semanticist can have no truck with primitives, since he must take the meanings of *all* words to be given by their inferential roles.⁴

I've heard the response that the inferential role semanticist should see the difference in meaning between 'cat' and 'dog' in terms of a difference between the typical looks of the two animals, in what it is like to see them, or, more behavioristically, in the capacity to recognize them. But it would be hard to justify including the look of a dog in the inferential role of 'dog' while rejecting the inference from 'dog' to 'barks' because it isn't analytic.⁵

Conclusion: The inferential role semanticist has nothing to gain from appealing to (traditional) analyticity as a principle of distinguishing among inferential liaisons.

It should be noted that the point just made and many other points in this paper involve examples of natural kind terms, where traditional claims of analyticity are widely thought to be relatively weak. Here I follow Fodor's and LePore's lead. Perhaps they have chosen the domain in part for its unfavorability to analyticity—since they wish to use Quine on analyticity as a club with which to beat the inferential role semanticist. Whatever their reason, a theory of what meaning is ought to be able to handle natural kind terms.

Now I will turn to the main point against the Crack. Suppose that we include in the inferential role of 'snake' inferences such as that from 'snake' to 'dangerous if rattling' and from 'rattling snake' to 'dangerous'. If we could include such inferential liaisons, then the inferential role of 'rattling snake' would be a function of the inferential role of 'snake' (and hence a function of the inferential role of 'snake' and 'rattling' together), and the argument for the Crack would collapse, because compositionality would be satisfied. Remember that the Fodor–LePore argument depended on the claim that if we include all the inferences in which a sentence participates

⁴ This is all oversimple. First, inferential role semantics might be said to make every word a primitive. Second, one can imagine an inferential role semantics that appeals to different classes of inferences, giving a privileged status to inferences that correspond to traditional definitions.

⁵ What should the inferential role semanticist do about the paucity of analytic inferences? I will argue later that inferential role semantics should be thought of in terms of narrow meaning and narrow content, and that this option can be used to avoid the problem. But even without this move, there are two obvious options: one is to include all inferential liaisons in the inferential roles, and the other is to somehow extend inferential roles outside the head into the world.

in its inferential role (*i.e.* we accept holism), we get non-compositional inferential roles, so if Fodor and LePore can't prevent including the inference from 'snake' to 'dangerous if rattling' in 'snake's inferential role, their Crack will crumble.

So: can we put *all* inferences on the Official List of inferential roles for sentences? And how will that affect the inferential roles of words? Focusing on the latter question, it is important to realize that no one is considering any alternative to including in the inferential role of a word all the inferences involved in inferential roles of sentences which contain that word. Fodor and LePore don't mention any alternative, nor as far as I know is any alternative ever discussed in the literature. As I said at the outset, it is sentences that in the first instance *have* inferential roles, and the natural way of thinking of the inferential roles of words and other sub-sentential constituents of sentences is as constructions out of the inferential roles of sentences. Whatever inferences one decides to include in the inferential roles of sentences will be automatically carried through to the level of words.

Fodor and LePore argue that meanings are compositional; but inferential roles are not compositional; so meanings cannot be inferential roles. But the idea that 'compositionality is an embarrassment' (p. 334) for inferential role semantics is most peculiar. A genuine issue about compositionality arises, for example, with respect to externalist semantic theories, where we have an independent purchase on the reference or satisfaction conditions of words, and the truth conditions of sentences. There is a real issue of whether the latter is a function of the former. But in the case of inferential role, where we have no independent purchase on the inferential role of words, no such issue arises. In the context of inferential role semantics, talk of accepting or rejecting compositionality is hard to understand without some non-compositional proposal and some reason why any inferential role semanticist might accept it. The most obvious non-compositional proposal would be to have *two* Official Lists (or Principles for Constructing Lists) of inferences to be included in inferential roles, one for sentences, and another for words (or more Lists for other subsentential constituents). I can't think of a reason why an inferential role semanticist would want to do that, and Fodor and LePore don't even hint at anything to do with this issue. Indeed, Fodor and LePore make it appear that there is a 'compositionality problem' for inferential role semantics only by explicitly assuming that the meaning of 'brown' is the property of being brown, and the meaning of 'cow' is the property of being a cow (p. 334), which is a *non-inferential-role* account.

Clarity about compositionality requires clarity about the distinction between compositionality and holism. Holism, for inferential role semantics, is a matter of *how many* inferences are included in the Official List of inferential roles. If all inferences are included, we get holism. The example of including in the Official List the inference from 'This is a rattling snake' to 'This is dangerous' is meant to be a proxy for holism.

Compositionality, by contrast, has to do with our stipulation of what the inferential roles of words are to be, *given* that one has already decided on the Official List of inferential roles of sentences.

In sum, Fodor and LePore are right that there is an issue of whether to include the inference from 'rattling snake' to 'dangerous' in our Official List of inferential roles. But if we *do* decide to include it in the Official List for sentences, *of course* we should include it in the Official List for words. So the real issue is one of holism, not compositionality.

Before I get to their argument, let me make some further points about compositionality. If holism is true, the semantic identity of each word in the language carries information concerning the semantic properties of all the words in the language. Holistic compositionality is hyper-compositionality. It would be natural to respond by framing a definition of 'compositionality' that would include the non-holistic compositionality that Fodor and LePore have in mind, while excluding hyper-compositionality, in which the meanings of each of 'snake', 'dangerous' and 'rattling' contain the information that rattling snakes are dangerous. But though this would not be hard, it would be pointless, since hyper-compositionality is genuine compositionality, though an extreme form of it. To redefine 'compositionality' to exclude hyper-compositionality, and then complain that inferential role theories are not compositional would be an ad hoc verbal maneuver of no interest.

We can clarify hyper-compositionality further by giving a more explicit proposal about the semantic values of words. First, let the semantic value of a sentence be an ordered pair whose first member is the set of inferences *to* that sentence and whose second member is the set of inferences *from* that sentence. (On many views of what is to be included in inferential roles, these will be infinite sets.) Then the 'semantic value' (in the sense of the semantic property the theory assigns) of a word could be taken to be the set of semantic values of sentences containing that word, *i.e.* a set of ordered pairs, one for each sentence containing the word.

These semantic values are compositional in the sense used here, *semantic* compositionality, namely there is a function from the semantic values of parts to the semantic values of wholes. (This is the sense that is used in formal semantics, and that appears in Fodor's and Lepore's Glossary (1992).)⁶

There is another sense of 'compositionality', *psychosemantic* compositionality, in which hyper-compositionality clearly does not involve compositionality. For a theory to be compositional in this sense is for it to tell us how a person *represents* the meanings of words and uses them to

⁶ The function will also be computable: If one wants to 'compute' the semantic value of 'Water floats ships' from the semantic value of 'water', one can search the ordered pairs in the semantic value for 'water' until one finds one in which 'Water floats ships' is repeated (in one of the sets in the pair): That ordered pair is the semantic value of the sentence.

form representations of the meanings of sentences. A holistic IRS does not do this *by itself* (nor do other non-psychosemantic theories, possible worlds semantics theories, for example), so it requires supplementation by a psycho-semantic theory. It is psychosemantic compositionality that is needed to explain such psychological phenomena as systematicity and productivity. As Fodor (1987) points out, it seems to be a fact about humans that if they can frame the thought that Mary loves John, then they can frame the thought that John loves Mary. A theory of how people represent the semantic and non-semantic properties of the 'parts' of these thoughts has the potential to explain such facts.

To the extent that a hyper-compositional theory seems wildly counter-intuitive, this in my experience is due to a confusion of these two senses of 'compositional'.

Hyper-compositionality is dramatically illustrated by attention to the probabilistic inferential role theory that Fodor and LePore mention. According to this theory, the inferential role of a sentence is a matter of its conditional probability relative to *every* other sentence.⁷ It is this conditionalizing on every other sentence that brings in the holism. Fodor and LePore argue against the compositionality of these probabilistic inferential roles as follows. They imagine that the (subjective) probability that a rattling snake is dangerous is high, whereas the probability that a snake is dangerous is low (most snakes are harmless), as is the probability that a rattling thing is dangerous. (Actually, they put this in terms of their brown cow example.) They argue that since the first is high, whereas the second and third are low, the semantics of the first can't be a function of the semantics of the second and third. There is some danger in thinking of compositionality in terms of relations among sentences—as opposed to the relations between sentences and subsentential constituents. But going along with this fiction, let's suppose that compositionality requires that the semantics of (1) be a function of the semantics of (2) and (3).

1. This rattling snake is dangerous.
2. This rattling thing is dangerous.
3. This snake is dangerous.

Now Fodor and LePore see the issue as settled by the fact that (1) has a high probability, whereas (2) and (3) have low probabilities. Without pausing to consider why a compositionality function can't take two low probabilities into a high probability, we can see an important problem. The probabilistic semantics for these sentences is a matter of their *conditional* probability relative to every other sentence. This semantic entity is a

⁷ I should mention that Hartry Field, in the article that Fodor and LePore mention, says he is giving a notion of sameness of meaning, not meaning itself. Also, since these are subjective probabilities, the criterion of sameness of meaning is intrapersonal, not interpersonal.

function from sentences to probabilities, not a simple probability. This function in the case of (2) includes (2)'s high probability conditional on 'This is a snake'. Likewise, this semantic function in the case of (3) includes its high probability relative to 'This is rattling'. In short, the probabilistic meaning function for (2) and (3) both include the information that rattling snakes are dangerous.

Leaving the topic of probabilistic inferential role, let us return to a main question of this section: just what do Fodor and LePore have to say against including the inference from 'rattling snake' to 'dangerous' in the inferential role for 'snake'? Recall that without a way of precluding this option, the argument that including all inferential liaisons in inferential roles yields non-compositional inferential roles will fail, and they will not be able to conclude that the inferential role semanticist has no choice but to accept analyticity. So their whole argument depends on this step.

No argument is given in (1991), but there is one in (1992): '... there is a way to reconcile analyticity with holism—namely by accepting a semantics that represents every inference as analytic ... the idea that all inferences are compositional is the idea that all inferences are analytic, and the idea that all inferences are analytic, though it, too, is compatible with holism, is, on independent grounds, perfectly mad.' (pp. 181–2) So the claim is that if we include the inference from 'rattling snake' to 'dangerous' in the inferential role for 'snake', we make 'Rattling snakes are dangerous' analytic, (or the inference from 'rattling snake' to 'dangerous' analytic) and that is bad.

Let us postpone the issue of *why* including the inference is supposed to make 'Rattling snakes are dangerous' analytic, and ask first why, *if* this sentence is analytic, that is supposed to be bad?

The sum total of the argument in (1992) is that this idea is 'preposterous on the face of it' (p. 164), 'patently preposterous' (p. 174), that it is not 'possible to take seriously' (p. 174), that 'surely this is preposterous' (p. 182), that is 'perfectly mad' (p. 182), and 'incapable of being taken seriously' (p. 183), and that it is '... an option which Quine, quite sensibly, didn't even bother to consider ...' (p. 183). No other argument appears in either publication.⁸ I have three points to make about all this. First, this sort of confidence about what is and is not analytic does not fit with their attempt to beat the inferential role semanticist about the head and shoulders with the club of Quinean rejection of analyticity. I will come back to this point. Second, is all the name-calling in part a way of invoking the Old Hat? If 'Rattling snakes are dangerous' is analytic in my idiolect but not in yours, then, they might say, how could it be that my accepting this sentence and your rejecting it amount to a disagreement? Adding

⁸ Well, actually there is an argument that has the role of an aside (1992, p. 182) to the effect that this sort of holism would preclude to the Quine/Duhem thesis. I won't discuss this point further.

analyticity to the Old Hat in the 'no analyticity' version presented at the outset of this paper perhaps makes it more dramatic, but I don't see that it adds any substance. Finally, philosophers are wont to take their intuitions about concepts very seriously. And there is some justification for this. Our concepts of consciousness, agency, belief and so on play an important role in our lives, and we are interested in an elucidation of them, warts and all. But analyticity is a technical concept whose warts are of only technical interest. In some arenas, what we want by way of theory is a 'good fit' to intuitions, but this methodology is not always appropriate. In semantics, what we want is a theory with some power and interest. If we have to get rid of intuitions about 'analyticity', *etc.*, so be it. (I would say the same about 'meaning', but this is more controversial.) These intuitions are largely the product of the discipline itself.

Katz (forthcoming) argues that Kant and Locke thought of analyticity as a narrow meaning containment relation. Frege introduced the current use in which truth plays such an important role. Thus our intuitions concerning 'analyticity' may vary with prevailing theory.

How might Fodor and LePore back up the claim that it is preposterous to regard the inference from 'rattling snake' to 'dangerous' as analytic? One possible Fodor–LePore answer would be that it is preposterous because if we were to regard all inferences as analytic we would not be able to recognize the existence of disagreement, change of mind, *etc.* If they were to give this answer, the Crack would be little more than the Old Hat.

What other answer could they give? Others might say that 'Rattling snakes are dangerous' isn't analytic or that there is no matter of fact as to whether it is analytic because nothing is analytic or because there are no facts of analyticity, or because analyticity is a hopeless confusion.⁹ Interestingly, Fodor and LePore have closed off these avenues of response. They endorse the analyticity of 'quasi-syntactic' inferences like that from 'brown cow' to 'brown', or to 'cow'. (I call these 'quasi-syntactic' because they depend on the repetition of a word.) The Quinean position that has been so influential rejects the very notion of analyticity as a hopeless confusion.¹⁰ By accepting *some* analyticities, Fodor and LePore pull the Quinean teeth from their critique.

What I am leading up to is this: Given that they accept *some* analyticities, how can they reject the analyticity of 'Rattling snakes are dangerous' without saying something like 'Well, *in fact* "dangerous" just isn't part of the meaning of "rattling snake"? Since the issue of whether 'dangerous' is or is not part of the meaning of 'rattling snake' is part of the very issue at hand (because we are discussing their argument against holism, *i.e.* against including all inferential liaisons in inferential roles), this plainly

⁹ Boghossian (forthcoming) distinguishes among a number of such anti-analyticity theses.

¹⁰ Horwich (1992, forthcoming) argues that Quine does not reject analyticity as a hopeless confusion.

would be begging the question. They are up to their necks in analyticity, so the issue about whether all inferences are to be included in inferential roles is just the issue of holism itself. The cash value of their argument is just the bare claim that holism itself is preposterous, and this claim has to be argued for, not just asserted. This is half of my response to their objection.

Well, what if Fodor and LePore were to try to evade the objection just given by becoming Quineans, rejecting analyticity as a confusion? This, they cannot do. Recall that I have been postponing the issue of why including the inference from 'rattling snake' to 'dangerous' in the inferential role for 'snake' is supposed to make 'Rattling snakes are dangerous' analytic, concentrating instead on why, if this sentence is analytic, that would be bad. But of course the two questions are closely related. Were Fodor and LePore to give a Quinean response to the latter question, their approach to the first would dissolve. That is, rejecting analyticity as a confusion would preclude their claiming that the holist has to regard 'Rattling snakes are dangerous' as analytic. This is obvious enough, but the points to be made in the next section will articulate it.

Further, their views on compositionality and its relation to analyticity would make it difficult for them to reject analyticity. '... it's plausible that compositionality entails analyticity *whether or not* you accept New Testament Semantics 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.' (p. 338)

I've alluded to the fact that Fodor and LePore tend to equivocate on 'analyticity'. Their official notion is truth in virtue of meaning.¹¹ But sometimes they seem to have in mind a distinct sense of 'analyticity', namely what one might call 'traditional analyticity', in which the extension of the term is tied to traditional examples, especially lexical cases like 'Cats are animals'. The point I've been making involves the official notion of analyticity. The holist whom they are arguing against says that the inference from 'rattling snake' to 'dangerous' is to be included in the meaning of 'snake'. To deny this on the ground that the inference is not in fact analytic is just to beg the question, if analyticity = truth in virtue of meaning. But what if the extension of 'analyticity' is pegged to traditional examples? This would block their argument in another way, for their argument starts with the claim that to include the inference from 'rattling

¹¹ They say (p. 332) '... the analytic/synthetic distinction is unprincipled. This means ... that there aren't any expressions that are true or false solely in virtue of what they mean.' Of course, the *true in virtue of meaning* analysis of analyticity applies to statements, not inferences. The corresponding idea for inferences would be validity in virtue of meaning. To avoid the complication of discussing analytic statements and inferences separately, I will stick with the formulation in terms of truth, ignoring validity.

snake' to 'dangerous' in the meaning of 'snake' would be to make this inference (and the sentence 'Rattling snakes are dangerous') *analytic*. If analyticity is restricted to traditional examples, this claim would certainly be false.

In sum, the Crack rests on the argument that if we include the inference from 'rattling snake' to 'dangerous' in the inferential role of 'snake', we make this inference (and 'Rattling snakes are dangerous') analytic, and that is bad. (Of course, this example is just a proxy for talk of including *all* inferences in inferential roles.) Here I have dealt with the latter, the *bad* part; the next section is about the former, the *analyticity* part.

What was my response to the *bad* part? I said that their rejection of the analyticity of 'Rattling snakes are dangerous' without argument amounts to the rejection of holism, without argument. But do I think that 'Rattling snakes are dangerous' is analytic or not? My answer, given below, is that inferential role semantics of the sort I favor is committed to the narrow meaning of 'dangerous' being included in the narrow meaning of 'rattling snake', so 'Rattling snakes are dangerous' has the narrow analog of analyticity. This involves no commitment to the analyticity of 'Rattling snakes are dangerous', as I will show below.

4. The Plausible-Sounding Principle and Compositionality

Let us now turn to Fodor's and LePore's argument that including the inference from 'rattling snake' to 'dangerous' in the inferential role of 'snake' makes this inference analytic. It is not too difficult to see what the argument is supposed to be. Fodor and LePore say: '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 its warrant to be determined by the inferential roles of its constituents.' So if the inference from 'rattling snake' to 'dangerous' is part of the inferential role and therefore the meaning of 'snake' (as the holist inferential semanticist claims), then this is an analytic inference. (This reasoning appears in the voice of an interlocutor on p. 337, but reappears in the authors' voice on p. 338.)

The argument presupposes something like the following principle:

The Plausible-Sounding Principle

Inferences that constitute an expression's meaning what it does are analytic.¹²

¹² Note, incidentally, that the Plausible-Sounding Principle cannot use 'analytic' as restricted to traditional sorts of examples, for then it would be of no use in arguing that the holist inferential role semanticist is committed to the analyticity of 'Rattling snakes are dangerous', since this isn't of the type of the traditional examples.

Fodor and LePore see the Plausible-Sounding Principle as a consequence of their view that, for inferences, 'compositionality entails analyticity' (p. 337).¹³ The idea is that '... if "brown cow" → "dangerous" is *compositional* [i.e. if this inference is part of the inferential role of 'brown' or 'cow'], it follows that it is also *analytic* ... the cost of representing an inference as compositional is that you then have to represent it as analytic' (1992, p. 181–2; italics in original). So if an inference is included in the inferential role of a sentence (such as the inference from 'This is a rattling snake' to 'This is dangerous'), and therefore is included in the inferential role of the constituent words (e.g. is part of the inferential role of 'snake'), then this inference is analytic. Hence the Plausible-Sounding Principle.

The rest of this paper is about the Plausible-Sounding Principle. Though Fodor and LePore don't make much of it, I see it as the corner-stone of the general line of thought that they are running to the effect that inferential role semantics runs into trouble with analyticity. Should they be tempted to pull back from a commitment to it, they would have to show how, without it, they can derive the analyticity of the inference from 'rattling snakes' to 'dangerous' or of 'Rattling snakes are dangerous' from the claim (of holistic inferential role semantics) that the inference from 'rattling snake' to 'dangerous' is to be included in the inferential role of 'snake'.

Before I get into a critical examination of the Plausible-Sounding Principle, I have a few remarks about analyticity. You will recall that Fodor's and LePore's official notion of analyticity is truth in virtue of meaning. As I mentioned, they also use 'analyticity' so as to restrict it to traditional examples. Now if it is analyticity in the latter sense which they are trying to foist on the inferential role semanticist, they *never* had an argument that the inferential role semanticist is committed to analyticity in the first place.

To see this, suppose that Fodor and LePore are right that it is unacceptable to regard 'dangerous if rattling' as part of the meaning of 'snake', and thus that the inferential role semanticist is unable to use 'All of them' as a principle of inclusion of inferential liaisons in inferential roles. (This is the step in the argument that I have been worrying about up until now.) They go on to claim that once holism is rejected, the inferential role semanticist has nothing left to fall back on as a principle of inclusion except analyticity. Now if analyticity here means 'true in virtue of meaning', or even 'reducible to logical truth via substitution of definitions', the argument has a certain *prima facie* appeal (though I will be arguing that there is an unnoticed assumption here that should be rejected). The idea would be that all and only inferences that the inferential role semanticist can independently justify as analytic should be included. But if 'analytic'

¹³ See the fuller quotation five paragraphs back. They go further, saying '... for inferences, analyticity and compositionality are the same thing ...' (1992, p. 185; italics in original).

involves a commitment to the *extension* of 'analytic' as traditionally viewed (including, for example, to 'Cats are animals'), then the idea that the inferential role semanticist has *nothing else to fall back on* seems more than strained. It is hard to fault the claim that the inferential role semanticist should have a *principle* of distinguishing between inferential liaisons that are in from those that are out. But it is another *more doubtful* thing to say that the only such principle is analyticity. (I will be arguing that analyticity won't do.) And it is still another *much more doubtful* thing to say further that the only option is to accept traditional ideas of what is included and what is excluded from that category, especially since as I mentioned earlier, there aren't enough of these traditional analyticities. Any holist who accepts the criterion of analyticity is committed to regarding 'Rattling snakes are dangerous' as analytic, and thus this holist is committed to rejecting the traditional extension of 'analytic'.

In sum, if Fodor and LePore stick to their official sense of 'analytic', then their emphasis on compositionality is misplaced, for the Plausible-Sounding Principle that they have to buy into anyway would, if true, require the inferential role semanticist to accept the analyticity of whatever inferences he includes in inferential roles. But if they use 'analytic' so as to involve a traditional idea of the extension of 'analytic', then their argument that the inferential role semanticist is committed to analyticity loses its plausibility.

Returning to the Plausible-Sounding Principle, I should add that the Fodor-LePore argument based on it is a special case of a more general argument developed in Boghossian (forthcoming), which I object to on the same ground as I object to the Plausible-Sounding Principle.¹⁴ Boghossian notes that philosophers have accepted Quine's argument against the analytic/synthetic distinction, but not Quine's attack on meaning itself. He argues that anyone who accepts determinate facts about what means what is committed to analyticity, so one has either to embrace both of the Quinian doctrines or neither of them.

The Plausible-Sounding Principle holds that for the special case of inferential role semantics, determinate meaning facts engender analyticity, but the considerations that favor the Plausible-Sounding Principle are general, and would apply to other theories of meaning as well (see the second paragraph of Section 3 for partial acknowledgement of this point). But if Fodor and LePore accept the general form of this argument, they are in trouble, for they themselves accept determinate meaning facts as part of their Old Testament semantics.¹⁵

I will give an argument against the Plausible Sounding Principle. It will

¹⁴ I accept Boghossian's claim for *wide* meaning, but not for narrow meaning. See below. Some of the points here made against Fodor and LePore were arrived at independently by Boghossian.

¹⁵ Bob Stalnaker made essentially the same point in comments sent to Fodor and LePore.

show how an inferential role semanticist can have holism plus compositionality without analyticity, and without succumbing to the Old Hat. It will also apply against Boghossian's argument. My point is very simple. It is that inferential role theories are much more plausibly taken as theories of narrow meaning than of wide meaning, and—this is the punch line—determinate meaning facts about *narrow* meaning do not engender analyticities. So the Plausible Sounding Principle (that inferences that constitute meaning are analytic) is wrong—for narrow meaning. This point also blocks the Old Hat, the argument discussed at the outset that holistic inferential role semantics cannot acknowledge agreement and disagreement.

It may be said that Fodor and LePore needn't be seen as arguing against narrow inferential role theories. Perhaps they are best taken as refuting a wide inferential role semantics in which the roles are seen as extending into the world? (This might be said, but actually their (1992) is explicit about being directed against both narrow and wide inferential role semantics. Indeed, they are at pains to point out that their argument doesn't depend on assuming that inferential role semantics is only an account of narrow content. See p. 168.) So how can I defend wide inferential role semantics against their attack by defending narrow inferential role semantics? It suffices to say that if their attack can be shown not to apply to the most plausible version of inferential role semantics, why should those who favor inferential role semantics be disturbed by it? The ball is (or soon will be) in their court on this issue.¹⁶

I should mention at this point that I am not a True Believer in inferential role semantics or in narrow meaning or content. Rather, I think both are interesting and promising theoretical options that will stand or fall together, and that no one has succeeded in working out successfully. Though I am not a True Believer, I do think that inferential role semantics is not

¹⁶ See Block (1987) for a discussion of the distinction between narrow and wide inferential role semantics. I discuss in that paper a 'two factor' inferential role semantics in which one factor is 'short-arm' inferential role that stops at the skin. This short-arm inferential role is meant to capture narrow meaning or content. The other factor is a referential/truth-conditional factor. I argue that this sort of theory is equivalent to a 'one-factor' theory of the sort favored by Harman (1982), in which the one 'long-arm' factor extends into the world. The holism issue reveals a difference between two different inferential role theories of wide content, though the difference doesn't line up exactly with the distinction between one factor and two factor theories. The two factor theory is equivalent to the one factor theory only if they are non-holistic, so my (1987) was mistaken. On the two factor theory, the holist has the option of making only the narrow factor holistic—this is the view defended here, a view that avoids the Old Hat. The one factor theorist, by contrast, if he is also a holist, will most naturally treat his entire factor holistically, yielding a distinct theory that succumbs to the Old Hat. In sum, the two most natural holistic inferential role theories of wide content are (a) a two factor theory that treats only the short-arm roles holistically, and (b) a one factor theory that treats the long-arm roles holistically.

refuted by the arguments given by Fodor and LePore, and that is the position I see myself as defending here.

There are many differences between thought and talk, and between the content of thought and the meaning of talk. But in what follows, I will tend to ignore these differences.

5. Two Factors and Disagreement

Let's agree that people do disagree, change their minds, *etc.* Now there is one obvious version of inferential role semantics that allows us to have our cake and eat it too—have holism and also disagreement, changes of mind, *etc.*

The version I have in mind involves a 'two factor' theory, one that combines the Old and New Testament approaches.¹⁷ According to one way of thinking of the two factors, there is an internal content factor—narrow content—that is holistic. And there is an external content factor, the theory of which is to be filled in, perhaps, by Fodor's own externalist account. Some object to this sort of view because they have no stomach for narrow content. This is not the place for an attempt to defend narrow content. I will rest with an *ad hominem* argument, namely, how can Fodor, Mr. Narrow Content, himself, say this?

Here's how the two factor view avoids the Old Hat: In thinking about disagreement, agreement, change of mind, and such, we let the wide contents individuate the propositional attitude contents. Suppose, temporarily, that we think of wide contents as purely referential/truth-conditional. If I first assert, then, in an about-face, deny 'Toads cause warts', we can individuate the contents as the one factor externalist would do so, concluding that I deny what I earlier asserted, despite the changes in narrow contents. In other words, for some purposes of attitude attribution, the equivalence relations among contents would ignore the narrow contents.

Similar points about individuation of computations are familiar. We can type-identify computations by function computed, ignoring the specific algorithm deployed and its implementation, or by the algorithm, or by the implementation. We can type identify bathtubs architecturally—in terms of size, shape, weight, and decorative properties. Alternatively, we can type identify them economically, in terms of initial price, usable life, and thermal insulation, ignoring architectural features. Equivalence relations can be based on any of these features.

Martin Davies has suggested that I point out that essentially this idea

¹⁷ See Dretske, 1988, Chapter 6.4, especially Note 2 on p. 150, for an indication of why one who starts with a single factor externalist theory (of the Old Testament sort that Fodor has expounded and that Fodor and LePore seem to be pointing to) might find himself moving in the two factor direction. See Harman, 1982, Horwich, 1990, for one factor accounts. See also Devitt, 1990.

about individuation of content can be appreciated without any commitment to narrow content. Suppose that the way that you think about Hesperus differs from the way I think about Hesperus. Nonetheless, the word 'Hesperus' reasonably may be taken to have the same meaning in our mouths. If meanings were individuated as finely as thought contents, no two people would be likely to share meanings. So we individuate in a more coarse grained manner. But we will also want to maintain that 'Hesperus' differs in meaning from 'Phosphorus'. The point is that we can think of thought contents as determining reference and truth conditions, but still being much more fine grained than the truth conditions that they determine. And we can think of meanings as determining truth conditions while still being more fine grained than the truth conditions, but not as fine-grained as the thought contents. We have a many-one relation between thought contents and meanings, and a many-one relation between meanings and truth-conditions. Thus we can see three different equivalence relations on detailed thought contents. In thinking about agreement and disagreement using this apparatus, we would individuate in terms of meaning or in terms of truth conditions, depending on context, though some contexts might also elicit individuation in terms of fine-grained thought contents.

Adding in narrow content complicates this simple hierarchical structure, because narrow contents are both more and less coarse-grained than truth-conditional contents. My 'Hesperus'-thought has the same narrow content as my Twin's 'Hesperus'-thought, despite their truth conditions being different—if the star referred to as 'Hesperus' here is distinct from the one referred to on Twin Earth by the same name. Though narrow contents are in this respect coarser grained than truth-conditional content, narrow contents are also finer grained, since my 'Hesperus'-narrow content differs from both your 'Hesperus'-narrow content and my 'Phosphorus'-narrow content.

In the last four paragraphs, I supposed that wide contents are purely referential/truth-conditional. On this way of talking, the wide content of the belief that we express with 'Hesperus shines' would be the same as the wide content of the belief that we express with 'Phosphorus shines'. On this familiar view of wide content, wide content is entirely within the purview of the external factor of the two-factor theory. The cognitive difference between the beliefs expressed by the two sentences just mentioned could, on this view, be captured in terms of different cognitive relations to a single wide content. (See Fodor, 1990, for a view of this sort.) A distinct but also familiar terminology regards wide content as not entirely referential/truth-conditional. This view gives the beliefs expressed by 'Hesperus shines' and 'Phosphorus shines' different wide contents. On this view of wide content, narrow content contributes somewhat to the individuation of wide content. Specifically, the difference in the inferential roles of the words 'Hesperus' and 'Phosphorus' makes for a difference of wide content, despite the fact that the two words refer to the same thing.

The point depends on the different roles of the two words, but of course a single word can have two different roles as well.¹⁸ I don't want my reply to Fodor and LePore to depend on either view of how wide contents are to be individuated. I can exploit whatever principle of individuation is favored by those who dislike holism.

Of course, since narrow contents are holistic, whereas truth-conditional contents are not, narrow contents will be individuated much more finely (even though they will also be individuated more coarsely, as in the Twin Earth examples). So: many differences in narrow content will not make for differences in truth-conditional content. It is reasonable to ask what narrow content is *for* (and also what truth-conditional content is *for*). Answer: narrow content has a role in psychological explanation.¹⁹ For purposes of certain kinds of psychological explanation, narrow content differences matter despite the fact that they don't make for differences in truth-conditional content, as Loar (1987) points out. By contrast, truth-conditional content attribution is useful for communication and other contexts where information is important, and where psychological differences don't matter. If the word gets around that the announcer has said that Bigbucks won the Kentucky Derby, all concerned can pick up the information without attending to how the announcer conceives of these things. The announcer's narrow contents matter to predicting and explaining his behavior, but not to the information he has communicated. According to this picture, it is wide contents—or perhaps referential/truth-conditional contents, if these differ from wide contents—that are attributed with *that* clauses.²⁰ Narrow contents, by contrast, are not paired 1–1 with wide contents, and hence not 'straightforwardly' attributable. We can, however, speak of my narrow content for 'Water puts out fires', or yours. 'You can't, in absolute strictness, express narrow content; but as we've seen, there are ways of sneaking up on it.' (Fodor, 1987, p. 51)

I have mentioned three categories of content: narrow content, truth-conditional content and wide content. I also mentioned two distinct views of wide content. According to the first, wide content reduces to truth-conditional content. On this view, Oedipus both accepts and rejects the wide content expressed by 'Oedipus is married to his mother', but via different modes of access to this wide content. According to the second view of wide content, Oedipus accepts one wide content (expressed by 'I am married to Jocasta'), but rejects a distinct content (expressed by 'I am

¹⁸ Pierre visits London twice, once in a foul mood in drizzling rain, again, as if to a different place, in high spirits and brilliant sunshine; he assumes that there are two places in England called 'London', one sparkling, one dreary. See Kripke, 1979; Loar, 1987.

¹⁹ Note that I am *not* claiming that narrow content is the only content that has a role in psychological explanation. Explaining behavior in terms of narrow contents does not preclude *also* giving some explanations of behavior in terms of wide contents.

²⁰ I differ from Loar in *accepting* wide contents; Loar confines himself to talk of wide attributions.

married to Mom'). Narrow content has a home in psychological explanation, truth-conditional content has a home in communication. I do not take a view on the utility of wide content to the extent that it differs from truth-conditional content.

6. Analyticity, Finally

'Not so fast', Fodor and LePore might say. 'What about all those analyticities? The analytic truths involving narrow contents that you are committed to *preclude the option* of individuating contents in terms of truth conditions, given your holistic narrow contents. You can't just pretend your narrow contents aren't there, since analyticity has consequences for truth. Suppose I used to think that toads cause warts, but now, having changed my mind, I accept that toads don't cause warts. The externalist says that the content of my earlier "Toads cause warts" is the same as what I now deny with "Toads don't cause warts". But you *can't* say that, since your views about constitutivity of narrow contents commit you to the claim that both the original assertion and its later denial are analytically true. How could one content be such that its assertion is analytically true, and its denial is analytically true too? Analytic truths are true in virtue of meaning, and since the meanings of some of the words must have changed, so have the contents.'

The Fodor and LePore reply that I am imagining is that the putative analytic truths involving narrow contents preclude treating the truth-conditional content of 'Toads cause warts' as the same when asserted as when denied. In short, I am imagining Fodor and LePore to be making a last-gasp appeal to a version of the Old Hat. In doing this, they will be conceding that the Crack does not stand on its own, but that will be cold comfort to me, if I am defeated by the Old Hat. If their objection is right, it will disable my refutation of the Plausible-Sounding Principle and my reply to Boghossian too.

Here's how I avoid this objection. I point out 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.²¹

I think it will help, but only as a first step in discussing this matter, to

²¹ Note that narrow content inferences cannot be regarded as inferences in any sense that involves *truth*. (See Block, 1987, for more on this point.) But this is OK with Fodor and LePore, since they say that they are happy to accept a causal interpretation of 'inference'—so long as it is not an interpretation of *analytic* inference. Note also that the proposal for defining the inferences that are in inferential roles as the analytic ones is ridiculous if inferential role semantics is taken as a theory of narrow content, for there can't be any analytic narrow content inferences.

consider my utterance of 'Feynman was a genius'. It is true. But now consider my Twin on a Twin Earth where Twin Feynman is not a genius, but only a front man for a shy genius who feeds Feynman all his lines. This is done sufficiently smoothly that no causal signals from Twin-Feynman to *my* Twin have been any different from the causal signals from the real Feynman that have reached me. My Twin says 'Feynman was a genius', but what he says is false. My Twin is as much like me as you like. So we have the same narrow contents. Conclusion: thoughts and utterances with the same narrow content can be true in one environment, false in another. Narrow contents have no truth values, and hence no truth conditions.

This point is only a first step because it can reasonably be objected that all the Feynman story shows is that truth conditions are relative to context. Such narrow contents might still have an analyticity-like property, namely, yielding a truth (in virtue of meaning) *in all contexts*. Let us call this property *hyper-analyticity*.²² A hyper-analytic narrow content makes any wide content of which it is the narrow content analytic. The Feynman example helps to focus attention on the real issue, namely can statements of the inferences that are constitutive of inferential roles involve true premises and *false* conditions? Or, alternatively, must these narrow contents always be hyper-analytic? The issue is hard to discuss without agreement on which inferences *are* constitutive of inferential roles. However, I think I can sidestep important disagreements about what inferences are constitutive by considering the cases that Fodor and LePore regard as the inferential role semanticists' last refuge, traditional lexical analyticities, including those that involve natural kind terms, for example 'Cats are animals', or rather the inference from 'cat' to 'animal'. (You will recall that they hope to stick the inferential role semanticist with a commitment to such analyticities.)

Consider my Aunt Bubbles' utterance of 'Glass is a solid'. Unknown to my aunt, her utterance is *false*: glass is actually a supercooled liquid—it has the amorphous molecular structure of liquids, not the regular structure of solids. If you wait long enough, your drinking glass will 'melt' into a puddle. (Really!! Ask any chemist.) There is a Twin Earth on which there is a Twin-Bubbles whose word 'solid' is anchored overwhelmingly to solids (as here), but whose 'glass' picks out a genuine solid. (What she calls 'glass' is not genuine glass.) My aunt's 'Glass is a solid' is false, but my Twin's Aunt's utterance of the same words is true. So utterances with the same narrow content as this putative analytic truth can be both true and false: the narrow content is not hyper-analytic.

The example depends on the fact that our 'solid' can be anchored firmly

²² I will sometimes speak of the hyper-analyticity of a sentence or of a wide content or proposition, meaning the hyper-analyticity of the normally associated narrow content. The definition of 'hyper-analytic' that I gave is only a first shot. One obvious refinement—unnecessary for my purposes, is to restrict the definition to contexts in which the narrow content does yield some wide meaning or other.

to solids despite our misidentification of glass as a solid. Here is another such example in which the mistake occurs on Twin Earth rather than here. Consider a Twin earth in which 'water' is used as here to refer to H_2O , but where water is very rare, most of the substances referred to as 'liquid' being slippery granular solids (Stephen White's (1982) example). 'Water is a liquid' said by them is false, although the same narrow content attaches to a truth here on earth.

I have been arguing against a hyper-analyticity for narrow meaning which a narrow meaning has if it must always yield a true wide meaning in a context. At least some narrow meanings are not hyper-analytic, as the 'Water is a liquid' and 'Glass is a solid' examples illustrate. Though these examples don't prove that there aren't other hyper-analyticities, the central role in inference enjoyed by 'class-inclusion' inferences suggests that no inferential role semantics could be built around hyper-analyticity.

A different analog of analyticity for narrow meanings is what I have been calling *constitutivity*.²³ The inference from 'water' to 'liquid' is part of what is constitutive of the inferential role of 'water'. But constitutivity is not the kind of analog of analyticity that will do Fodor and LePore any good, since constitutive inferences can be false, as I pointed out above, and their (imagined) appeal to analyticity in the beginning of this section required a connection between analyticity and truth.

Earlier, I said I would argue against the Plausible-Sounding Principle that inferences that are included in (and therefore part of what is constitutive of) inferential roles are thereby analytic. And I said that the same argument would suffice to show the error of Fodor's and LePore's claim that compositional inferences are analytic, and Boghossian's claim that commitment to determinate meaning requires a commitment to determinate analyticity. The arguments just presented are intended to fill the bill. Narrow meanings don't determine any truth conditions at all, so Boghossian's argument is stopped short—at least for narrow meaning. And inferences among narrow meanings can't be analytic, so the Plausible-Sounding Principle is disabled and compositionality doesn't entail analyticity. I also considered an imagined reply involving the substitution of hyper-analyticity for analyticity in Fodor's and LePore's and Boghossian's argument, and I argued that this move would gain them nothing.

Fodor and LePore also argue that since the inferential role semanticist must accept analyticity as the criterion for inclusion in inferential roles, he must give up on naturalizing those roles, that is, he must give up on characterizing these roles causally (p. 336). Without suggesting optimism about the prospects for causal characterization of inferential roles, I want to point out that once we see that the inferential role semanticist can think of inferential roles narrowly, escaping any analyticity requirement, the

²³ See Katz, 1990, p. 192, and Katz, forthcoming, for an account of constitutivity (though he doesn't call it that) that applies outside the context of inferential role semantics.

Fodor–LePore objection to naturalization loses its force. They also argue (in 1992) that inferential role semantics, holism, compositionality and no analytic/synthetic distinction are an inconsistent quadruple. But this is not so for holism of *narrow* contents. Fodor and LePore also argue that the need for an analyticity requirement deprives the inferential role semanticist of his favorite argument for holism: ‘... the present situation is rife with ironies ... inferential role semantics doesn’t *have* any holistic implications after all ...’ (p. 339). But with the disappearance of the analyticity requirement, that argument also dissolves.

Of course, this line of reply carries with it a promissory note of gargantuan proportions—to fill in the sketchy remarks that are often made about inferential role theories. I don’t know if this promissory note will ever be cashed, but surely we should not believe arguments that depend on its uncashability—such as the Fodor–LePore and Boghossian arguments—without an argument that it is uncashable.

One final comment on the Boghossian argument: I am often asked why narrow meaning facts are supposed to be *meaning* facts and why narrow content facts are supposed to be *content* facts. As I mentioned, meaning and content have two different homes, one in the realm of communication, the other in the realm of psychological explanation. The claim of narrow meaning to be meaning rests mainly in the latter home. This is why theories that distinguish two components of meaning (e.g. sense and reference) typically appeal to examples that involve same reference/different cognitive significance. It is these differences in cognitive significance that make for differences in psychological explanation. Further, narrow contents and meaning determine truth conditions *relative to contexts*, so there is an intimate relation between narrow meaning and content on the one hand, and truth on the other. That’s enough for me.

Note that I have not disputed Boghossian’s argument as it applies to wide contents. So, given that I accept wide contents, if Boghossian is right, I will have to accept facts of the matter about analytic relations among these wide contents, so isn’t my reply to Boghossian peripheral to his purposes? Perhaps, but my reply to Boghossian nonetheless plays an important role in my dispute with Fodor and LePore. If Boghossian’s point were right about narrow content, then I could not defend holistic narrow contents without giving up agreement and disagreement. I do not defend holistic wide contents, so Boghossian’s argument, if it is right, does not give a boost to Fodor and LePore.

Boghossian has pointed out (in conversation) that my example of the narrow content of ‘Glass is a solid’ could be circumvented if we restricted attention to what might be called ‘quasi-hyper-analyticities’, narrow contents that generate truths in all contexts *in which they generate the same wide content*. This maneuver may have a point for some purpose, but not for the purpose of this paper, which is to show how a holistic inferential role semantics can avoid both the Old Hat and the Fodor–LePore argument. Boghossian has argued that a commitment to determinate meaning

requires a commitment to determinate analyticity. I have pointed out that a commitment to determinate narrow meaning does not require a commitment to determinate analyticity (or hyper-analyticity) involving these meanings, and this allows me to avoid the Fodor–LePore argument and the Old Hat. For these purposes, it does not matter whether a commitment to determinate meaning requires a commitment to determinate quasi-hyper-analyticity.

7. More Objections

Jerry Fodor responded to an earlier version of this paper by saying that ‘stereotypes’ provide an obvious counterexample to the Plausible Sounding Principle. My Aunt Bubbles and her Twin share the inference from something’s having the stereotype of glass to it having the stereotype of solid. Though one inference is correct, the other is not. So the reasoning that I gave about narrow contents goes through equally well with stereotypes.

The first thing to say is that Fodor and LePore *need* the Plausible Sounding Principle, so to the extent that its refutation is obvious, too bad for them. Note that it is no trick to refute the Plausible Sounding Principle (or Boghossian’s argument for that matter) by adopting a *plainly inadequate* theory of wide meaning. If wide meanings are held to be *mental images*, for example, both the Plausible Sounding Principle and Boghossian’s argument will fail. As Putnam (1975) noted, meanings in any sense in which they determine reference cannot be *e.g.* images or stereotypes, since Twins can have the same ‘water’-images, but different ‘water’-references. If meaning determines reference, then different references require different meanings; hence *meanings ain’t in the head*, as Putnam emphasizes, and the Twin Earth word ‘water’ cannot be translated by our ‘water’. Stereotypes and images can be immediately disqualified as *wide* meanings, though they can be taken as candidate *narrow* meanings, in which case the Fodor point under discussion reduces to my point about narrow meaning.

Jerry Fodor has also argued (in conversation) that my ‘Glass is a solid’ argument does not challenge the hyper-analytic status of what he regards as compositional analyticities such as ‘Brown cows are cows’.

To see what the point might be, let us go over the role of hyper-analyticity in the argument I gave. Fodor and LePore make a number of claims involving analyticity. I pointed out that a plausible inferential role semantics deals in narrow contents which are never true and therefore never analytic, and further, that narrow contents of the putative ‘lexical analyticities’ that Fodor and LePore try to stick the inferential role semanticist with are not always narrow contents of analyticities—that is, they are not hyper-analytic—so hyper-analyticity will not substitute for analyticity in the Fodor–LePore argument. How could the putative hyper-analyticity of ‘Brown cows are cows’ resuscitate the Fodor–LePore argument?

Recall how the Fodor–LePore argument goes: The inferential role seman-

ticist is not supposed to be able to include the inference from 'rattling snake' to 'dangerous' in the inferential role for 'snake' because, according to Fodor and LePore, that would make this inference analytic, which, they say, it is not. Then Fodor and LePore say that the inferential role semanticist has to fall back on analyticity as a criterion for inclusion of inferences in his inferential roles. So there are two distinct appeals to analyticity: (a) If all inferences are included in inferential roles, then all inferences would be analytic; but some are not. (b) Analyticity is the only fall-back criterion.

Can the putative analyticity of 'Brown cows are cows' save a version of either (a) or (b) with 'hyper-analyticity' substituted for 'analyticity'? The point against (b) is easy to see. These putative 'quasi-syntactic hyper-analyticities' are not something that any serious inferential role semantics could fall back on, because there aren't enough of them. Such an inferential role semantics could make such semantic distinctions such as that between 'brown' and 'fake', but not between 'brown' and 'blue'. If 'quasi-syntactic' inferences (in which a word is repeated) are the only hyper-analyticities, then an inferential role semantics that restricted itself to hyper-analyticities would be a non-starter.

As for (a), I pointed out earlier that the claim that the inference from 'rattling snake' to 'dangerous' is not *really* analytic is question-begging in the context of a dispute about holism—whether being dangerous if rattling is genuinely part of the meaning of 'snake'. The putative hyper-analyticity of 'Brown cows are cows' is irrelevant to this issue. Certainly it gives us no reason to think that including the inference from 'rattling snake' to 'dangerous' in the inferential role for 'snake' will make this inference hyper-analytic.

I also pointed out that the inference to analyticity depended on the Plausible Sounding Principle, which is disarmed by the 'Glass is a solid' argument. The Plausible Sounding Principle is thus falsified—even a 'hyper-analyticity' version of it—despite the possibility that there are *some* cases of inferences that are constitutive of inferential roles that are hyper-analytic. A principle is falsified by exceptions even if there are some cases that comport with it.

Finally, Burgean arguments can be brought to bear on the putative hyper-analyticity of 'Brown cows are cows'. Consider a Twin of me whose language community uses 'brown' to mean *fake*. His deference to his language community will lead *him* to regard his utterance of 'Brown cows are cows' as false (let us say) when he learns the true meaning of 'brown' in his language community. Shouldn't we agree? Objections may appeal to 'idiolect' meaning, but given my Twin's insistence that what he said was false rather than idiosyncratic and true, the idiolect move can be fended off in the same way that Burge fends off the 'tharthritis' reply.

A quite different sort of objection is suggested by parts of Chapter 5 of Fodor's and LePore's (1992), and by Fodor (1987), Chapter 3. I have the impression that Fodor and LePore think that the two-factor approach that I endorse is doomed for a reason not yet mentioned, and so they may feel

justified in letting their argument presuppose its falsity. Their objection to the two factor theory is that there is nothing to keep the two factors stuck together. Why, they ask, can't you have a sentence that has the inferential role of 'Water is wet', but is true iff 4 is a prime number? I admit to being puzzled about why this is supposed to be an objection to inferential role theory.²⁴ What puzzles me about the objection is that though there is an interesting issue here, I don't see why there is supposed to be more of a *problem* for the two factor theorist than for anyone else. The issue arises for anyone who accepts the *existence* of inferential roles and of truth conditions, whether or not they think there is any plausibility in an inferential role semantics. Fodor and LePore never suggest that either inferential roles or truth conditions don't *exist*; indeed, their arguments presuppose that inferential roles (and truth conditions) do exist, so the issue they raise is as much an issue for them as it is for the inferential role semanticist.

Addendum

Having seen a draft of Fodor's and Lepore's reply, I cannot resist adding a short note: You will recall the distinction mentioned above between semantic compositionality and psychosemantic compositionality. The former consists in the existence of a function from the semantic values of parts to wholes; the latter has to do with how people represent and process these semantic values, and thus it is the latter that can explain productivity and systematicity. (Systematicity is the fact that someone who can think the thought that John loves Mary can also think the thought that Mary loves John; this is a psychological fact, having to do with how human thought works and is not to be explained by a semantic theory; there are possible thinkers and speakers whose language has a semantics like ours but whose thought is not as systematic as ours.)

Inferential role semantics (like possible worlds semantics) is compositional in the former sense—it specifies the relevant function—though it has little to say about the latter. Now that it has become clear (at least to me) that it is the latter sense of 'compositionality' that Fodor and Lepore have in mind, the flaw in their main argument can be seen to be their claim that compositionality is non-negotiable. It is no responsibility of a purely semantic theory to say how people represent or process the meanings of elements of their language. Indeed, part of the role of a purely

²⁴ My response to their question is that I *don't know* whether there could be a Twin of mine whose thought or utterance of 'Water is wet' shared the inferential role of my thought or utterance of those words even though his embedding in his world was such that its truth conditions had to do with numbers. I suspect that if you run with Burge's arguments about deference to the language community, you could make a case for this possibility.

semantic theory is to provide a framework for posing questions about processing. (I am indebted to Bob Stalnaker here.) So, for example, on a possible worlds account, any candidate for how people represent predicates will have to determine a mapping of worlds onto extensions.

A last remark about Boghossian's response: just as narrow meanings cannot be true solely in virtue of meaning, for the same reason they cannot be asserted solely in virtue of meaning. To assert something is to commit oneself to its truth, and so what cannot be true cannot be (warrantedly) asserted at all. Further, narrow contents in the class at issue in my paper aren't even parts of wide contents that are assertible solely in virtue of meaning. The same argument that shows that 'water is a liquid' isn't hyper-analytic also shows that it isn't assertible solely in virtue of meaning. It is part of our commitment in the use of natural kind terms that the world has a role to play in determining the truth value of what we say, and this is reflected in our implicit recognition that the world can thwart our assertions.

*Department of Linguistics and Philosophy
Massachusetts Institute of Technology
Cambridge, MA 02138
USA*

References

- Block, N. 1986: Advertisement for a Semantics for Psychology. In P. French *et al* (eds.), *Midwest Studies in Philosophy Volume 10*. Minneapolis: University of Minnesota Press.
- Block, N. 1987: Functional Role and Truth Conditions. *Proceedings of the Aristotelian Society Supplementary Volume 61*, 157–81.
- Block, N. 1991: What Narrow Content is Not. In B. Loewer and G. Rey (eds.), *Fodor and His Critics*. Oxford: Basil Blackwell.
- Boghossian, P. forthcoming: Analyticity. In C. Wright and R. Hale (eds.), *Blackwell Companion to the Philosophy of Language*. Oxford: Blackwell Publishers.
- Devitt, M. 1990: A Narrow Representational Theory of Mind. In W. Lycan (ed.), *Mind and Cognition: A Reader*. Oxford: Basil Blackwell.
- Dretske, F. 1988: *Explaining Behavior*. Cambridge, MA.: MIT Press.
- Field, H. 1977: Logic, Meaning and Conceptual Role. *Journal of Philosophy*, 74, 379–409.
- Fodor, J. 1987: *Psychosemantics*. Cambridge, MA.: MIT Press.
- Fodor, J. 1990: Substitution Arguments and the Individuation of Beliefs. In G. Boolos (ed.), *Meaning and Method: Essays in Honor of Hilary Putnam*. Cambridge: Cambridge University Press, 63–78.
- 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: Basil Blackwell.

- Harman, G. 1982: Conceptual Role Semantics. *Notre Dame Journal of Formal Logic*, 23, 242–56.
- Horwich, P. 1990: *Truth*. Oxford: Basil Blackwell.
- Horwich, P. 1992: Quine versus Chomsky on the Analytic–Synthetic Distinction. *Proceedings of the Aristotelian Society*, 92, 95–108.
- Horwich, P. Forthcoming: Scientific Conceptions of Language and Their Philosophical Import.
- Katz, J. 1972: *Semantic Theory*. New York: Harper and Row.
- Katz, J. 1990: *The Metaphysics of Meaning*. Cambridge, MA.: MIT Press.
- Katz, J. Forthcoming: The New Intensionalism.
- Kripke, S. 1979: A Puzzle about Belief. In A. Margalit (ed.), *Meaning and Use*. Dordrecht: Reidel.
- Loar, B. 1987: Social Content and Psychological Content. In R. Grimm and D. Merrill (eds.), *Contents of Thought: Proceedings of the 1985 Oberlin Colloquium in Philosophy*. University of Arizona Press: Tucson.
- Putnam, H. 1975: The Meaning of ‘Meaning’. In *Mind, Language and Reality: Philosophical Papers, Volume 2*. Cambridge: Cambridge University Press.
- White, Stephen, 1982: Partial Character and the Language of Thought. *Pacific Philosophical Quarterly*, 63, 347–65. Reprinted in *The Unity of the Self*. Cambridge, MA.: MIT Press, 1951.