

## Do inferential roles compose?

Mark McCULLAGH<sup>†</sup>

### ABSTRACT

Jerry Fodor and Ernie Lepore have argued that inferential roles are not compositional. It is unclear, however, whether the theories at which they aim their objection are obliged to meet the strong compositionality requirement they have in mind. But even if that requirement is accepted, the data they adduce can in fact be derived from an inferential-role theory that meets it. I explain this in terms of Robert Brandom's substitutional conception of inferential roles. The several objections the proposal invites are worth discussing because they rest, I think, on neglect of some interesting and important facts about inferential roles. Whether Fodor's and Lepore's strong compositionality requirement is justified or not, then, inferential-role theories do not have the problem that they claim to have identified.

Jerry Fodor and Ernie Lepore have argued – “frequently, loudly, and in many places”<sup>1</sup> – that inferential roles are not compositional. If they are right then things look grim for any semantic theory that incorporates a notion of inferential role. One problem with their argument is that it is unclear whether the theories at which they aim their objection are obliged to meet the strong compositionality requirement they have in mind. But even if that requirement is accepted, the data they adduce can in fact be derived from an inferential-role theory that meets it. I explain this in terms of Robert Brandom's (1994) substitutional conception of inferential roles. The several objections the proposal invites are worth discussing because they rest, I think, on neglect of some interesting and important facts about inferential roles. Whether Fodor's and Lepore's strong compositionality requirement is justified or not, then, inferential-role theories do not have the problem that they claim to have identified.

### 1. The objection

Fodor and Lepore support their objection on examples such as the following.

Suppose... that you happen to think that brown cows are dangerous; then it's part of the inferential role of “brown cow” for you that it does (or can) figure in inferences of the form “brown cow → dangerous.” But, at first blush anyhow, this fact about the inferential role of “brown cow” doesn't seem to derive from facts about the inferential roles of its constituents in the way that, for example, the validity of inferences like “brown cow → brown animal” or “brown cow → not green cow”

<sup>†</sup> Dept of Philosophy, University of Guelph, Guelph ON Canada N1G 2W1. Email: mmcculla@uoguelph.ca.

<sup>1</sup> 2001, 472. See their 1991 and 1992; perhaps there are other presentations as well.

might plausibly be thought to do. “Brown cow” entails “brown animal” because “cow” entails “animal”; “brown cow” entails “not green cow” because “brown” entails “not green.” But it doesn’t look as if either “brown” or “cow” entails “dangerous,” so, to this extent, it doesn’t look as if the inference from “brown cow” to “dangerous” is compositional. (1992, 177)

They recognize that they cannot direct this objection against those who hold that an expression’s role comprises only “analytic” inferences (178-79). For in that case the simple rejoinder is that while the entailment from “is a cow” to “is an animal” is analytic, the entailment from “is a brown cow” to “is dangerous” is not – so its not being compositionally derivable is not a problem. The objection is directed, then, at those inferential-role theorists who wish their theories to cover many more inferential relations than just those that would merit the label “analytic.” This point will be relevant below.

What sort of compositionality do Fodor and Lepore have in mind? The weakest commitment that deserves the title is to the idea that a theory be finitely stateable. One could in addition be committed to the idea that the statements in such a finite theory concern only the syntactically atomic expressions in the language. Call the former *finitistic compositionality* and the latter, stronger condition *atomistic compositionality*.

The reason usually given (e.g. by Davidson 1965) for requiring finitistic compositionality in a theory of meaning is that a language for which there was no finitely stateable theory of meaning would seem to be one that no creature like us could learn, hence, no theory of meaning that violates finitistic compositionality could be of any use in explaining the possibility of our learning the meanings of the sentences we use. If one takes that to be part of what a theory of meaning should allow us to explain, one will have good reason to require finitistic compositionality.

Atomistic compositionality, a stronger condition, is motivated more by intuition than by theoretical requirement. The intuition is that it is by knowing the meanings of its component words that we know the meaning of a sentence.<sup>2</sup> That seems an unobjectionable claim.

Fodor and Lepore presuppose atomistic compositionality (1992, 175, following Block 1986, 616). Why? One reason might be that their example would be untroubling if one could claim that the entailment from “is a brown cow” to “is dangerous” is in the compositional base. That possibility is excluded by the atomistic compositionality requirement but not by the finitistic compositionality requirement. Another reason might be that they share the intuition that supports that requirement. But that is problematic in the context of their argument. For it is far from clear whether the apparent truth of that intuition imposes any theoretical constraint on the inferential-role theorists that Fodor and Lepore take as their target. Recall that these theorists are, by Fodor’s and Lepore’s own description, ones whose aim includes explaining such inferences as that from “is a brown cow” to “is dangerous.” Yet nobody would say that knowing that “George is a brown cow” entails “George is dangerous” is part of what we *ordinarily* mean by “knowing the meaning” of either of these sentences – and it is what we *ordinarily* mean by that that figures in the intuition that supports the atomistic compositionality requirement. It is much more plausible to say that the inferential relations involved in (what we ordinarily call) “knowing the meaning” of a sentence are the ones that would traditionally count as analytic. If the intuition that supports the atomistic compositionality requirement is about any inferences, it is about those ones. Yet as was noted above, they are

<sup>2</sup> Qualification is needed as concerns idiomatic phrases but (as is common) I shall ignore that.

just a proper subset of the inferential relations that interest those theorists at whom Fodor and Lepore aim their objection. As the intuition that supports the atomistic compositionality requirement is silent on this broader class of inferences, Fodor and Lepore cannot appeal to that intuition to justify imposing that requirement. (The point is not that it is *inconsistent* both to demand that an inferential-role theory explain more than just the analytic inferential relations and to demand that the explanations be atomistically compositional; it is that the latter demand is *unmotivated* in the context of the former.)

Accordingly, one way for the inferential-role theorist to avoid Fodor's and Lepore's objection is to deny that her theory must be atomistically compositional. In support of that move she may say that the intuition about atomistic compositionality concerns inferences that form a proper subset of the ones she wants to explain. More generally, such a theorist's claim would be that she means to replace, or precisify in one direction (among several possible), the ordinary intuitive notion of linguistic meaning – the notion that figures in the intuitions that support the atomistic compositionality requirement.<sup>3</sup> (Indeed, if Paul Boghossian is correct to claim that nobody who rejects the analytic-synthetic distinction “can plausibly stop short of a radical indeterminacy about meaning” (1996, 30), then *all* inferentialists who are targets of Fodor's and Lepore's objection should be construed as having such a revisionary aim.)

It is far from clear, then, whether an inferential-role theorist is obliged to respect atomistic compositionality, given that the intuition supporting that requirement is expressed in terms of a concept she is, perhaps, meaning to improve upon rather than give a theory of. Rejecting the requirement of atomistic compositionality would allow the inferential-role theorist to put the “brown cow → dangerous” inference-form in the compositional base and to declare herself satisfied with merely finitistic compositionality, which Fodor's and Lepore's example gives us no reason to worry about. (I defend that last claim in a moment.)

## 2. Explaining the data atomistically

But even the inferential-role theorist who accedes to Fodor's and Lepore's demand for atomistic compositionality<sup>4</sup> can reply to their objection. For contrary to what they claim, it is possible to handle the relevant cases solely in terms of the inferential roles of syntactic atoms.<sup>5</sup> Technically it is trivial, but some interesting questions about inferential roles arise in the course of dealing with some objections that it invites due to its odd-

<sup>3</sup> Block qualifies his uses of the word “meaning” in describing his “two factor” proposal: “Nothing in my position requires me to regard narrow meaning and narrow content as (respectively) *kinds* of meaning and content.... I regard them as aspects of or as *determinants* of meaning and content” (1986, 626). Brandom's is a more radical departure from the traditional notion. He writes that one can “think... of communication in terms of sharing a relation to one and the same *thing* (grasping a common meaning) or in terms of *cooperating* in a joint *activity* (coordinating social perspectives by keeping deontic score according to common practices)” (1994, 479). Brandom rejects the former, traditional conception in favour of the latter, according to which “What is shared [by speakers engaged in communication] is a capacity to navigate and traverse differences in points of view, to specify contents from different points of view” (485).

<sup>4</sup> Unnecessarily, by my lights, Robert Brandom appears to do so (1994, 366, 374).

<sup>5</sup> Fodor and Lepore briefly seem to attempt doing this, but halfheartedly (1992, 178). For the only inferences they consider assigning to “brown” and “cow” are ones that traditionally count as analytic – “brown” to “not-green,” “cow” to “animal” – which is odd given their recognition that their objection applies only to theories that do not aim to explain all and only inferences of that sort. (It is worth noting, moreover, that those examples are intrinsically problematic since it is far from obvious precisely how to characterize their instances – see below, n. 9.)

ness. I will make the point in terms of Robert Brandom's account simply because it is a recent and well-worked-out inferential-role theory; we should expect to be able to make the point in terms of any equally worked-out inferential-role theory.

Brandom conceives of the inferential roles of non-logical vocabulary in terms of "simple material substitution-inferential commitments," or SMSICs. A SMSIC relating two syntactically simple expressions  $e$  and  $f$  determines two sets of sentence-frames (functions from expressions to sentences):<sup>6</sup>

$A = \{\text{sentence-frames such that inferences from } e\text{- to } f\text{-instances are proper}\}$

$B = \{\text{sentence-frames such that inferences from } f\text{- to } e\text{-instances are proper}\}$

This apparatus can handle Fodor's and Lepore's example as follows. We say that there is indeed a SMSIC determining a relation between "brown" and "dangerous": the members of its  $A$  set are sentence-frames such as "Bertha is  $\alpha$  and is a cow," "The thing in the pen is  $\alpha$  and is a cow," and so on. (Its  $B$  set includes sentence-frames such as "The thing in the pen is a cow but is not  $\alpha$ .") The idea is that an inference from a "brown" instance of a frame to a "dangerous" instance is permitted just in case that frame is one in which the expression replacing " $\alpha$ " occurs in conjunction with "is a cow." What we are doing is encoding as a property of the syntactically atomic expression "brown" something that is more naturally thought of as a property of the complex expression "brown cow." Because this is possible, there is at least one notion of inferential role that allows us to say that "brown cow  $\rightarrow$  dangerous" inferences *are* part of the inferential role of "brown."

Two preliminary points should be made about this proposal. First, it is not enough for it to be atomistic, as it obviously is; to count as atomistically compositional, the proposal must not require that there be infinitely many such clauses in a full description of a language. Infinitely many such clauses would be required if there were infinitely many examples such as Fodor's and Lepore's. Their example is of a complex expression ("brown cow") whose inferential role (entailing "is dangerous") is determined by a belief (that brown cows are dangerous). For there to be infinitely many such cases there would have to be infinitely many such beliefs. Now, although there is a sense in which each of us has infinitely many beliefs, all but finitely many of these are derivable by inference from the finitely many beliefs that our brains store at any given moment, and inferentially derived beliefs obviously do not give rise to worries about the non-compositionality of inferential roles. So there could be only finitely many cases of the "brown cow  $\rightarrow$  dangerous" sort – thus, only finitely many cases that would call for the treatment I have described.

Second, the proposal is unnecessary for languages with logical vocabulary. For in such a case the speaker will believe that if something is a cow and is brown then it is dangerous, in which case the inferential roles of "if... then" and "and" will allow for the derivation of "brown cows are dangerous" without the need for this inference to be built in, as it were, to the role of "brown." Fodor's and Lepore's example is not even *prima facie* troubling unless we are trying to derive the "brown cow  $\rightarrow$  dangerous" inferences by appealing *solely* to inferential roles assigned to non-logical vocabulary. (This may explain why inferentialists have not discussed examples like the ones Fodor and Lepore base their objection on. Inferentialists are fond of using the inferential roles of logical connectives to illustrate what inferential roles are – clearly they envision their theories applying to languages with such vocabulary.)<sup>7</sup>

<sup>6</sup> See Brandom 1994, 373-74.

<sup>7</sup> Block, for instance, discusses only the inferential role of "and" when explicitly addressing the compositionality of inferential roles (1986, 667).

Now to the objections. Probably what will have come to the reader's mind is, "Huh? You call that an inferential role? It's bizarre and *ad hoc*." I think there are several components to this objection and I hope to address it by treating them individually.

One objection is that we should not allow inferential roles to be of such restricted generality. No theorist who speaks of inferential roles has in mind a conception according to which, say, "blig" entails "blog" only on Wednesdays. It could strike someone as similarly illegitimate to hold that "brown" entails "dangerous" only when the former occurs in conjunction with "is a cow." This objection rests on the intuition that the proposal violates some basic, simple principle governing the generality of inferential roles. If there were some such principle, it might be a fine objection. But there is not. Even the inferential relation between "brown" and "coloured" – despite being about as straightforward as they come – is restricted in its generality: you can infer from "brown" to "coloured" in sentences such as "John's suit is brown" but not in sentences such as "Nothing in Gerald's room is brown," "If Gerald's suit is brown, it is not red," "Gerald will attack anything that is brown" or "Gerald said that the paint is brown."<sup>8</sup> Almost all inferential relations are of restricted generality.<sup>9</sup> So it is no objection to the relation just proposed between "brown" and "dangerous," that it is of restricted generality.

Another objection is that the generality of the specified inferential relation between "brown" and "dangerous" differs from the generality of inferential relations between other expressions. It does. But it is hard to see why that is objectionable. Indeed, Brandom argues at some length (1994, ch. 6) that singular terms and predicates have inferential roles of different generality. (He argues that the former are symmetric – the *A*- and *B*-sets are the same – while the latter are asymmetric.) And Fodor and Lepore themselves appeal (in an objection to Brandom's criterion of singular termhood) to examples of such differences even among expressions in the same syntactic category. (They claim that "Oxford" is substitutable for "Magdalen" in the frame "Father was at  $\alpha$ " and that "his wrist" is substitutable for "his whole arm" in the frame "John painted  $\alpha$  red" (2001, 476). But obviously the former substitution doesn't work in the latter frame, and *vice versa*.)

Another objection is to the *ad hoc*-ness of the proposal. Why build the relation to "dangerous" into the inferential role of "brown"? Recall the second point I made about the proposal: that it is unnecessary in languages with logical vocabulary. These are bizarrely impoverished speakers we are theorizing about; of course the theory is bizarre. And the motivation for the proposal is entirely legitimate: to make it possible to derive the inferential roles of syntactic complexes from the inferential roles of syntactic simples. That is the restriction that Fodor and Lepore impose; it is hardly fair to label a straightforward attempt to satisfy it *ad hoc*. (Indeed it is the restriction, rather than the footwork required to satisfy it, that deserves the label.)

Nor is it a successful objection to say that the specified inferential relation is of unprojectible generality. All that is being credited to speakers, when we discern this inferential role in their discourse, is an ability to discern, *among* those occurrences in which "is brown  $\rightarrow$  is coloured" inferences are proper (which we might call the "stan-

<sup>8</sup> See Brandom 1994, 398-99 for a discussion of this issue.

<sup>9</sup> This makes necessary a qualification on Fodor's and Lepore's use (and my use for the purpose of discussing their objection) of the locution inferences of the form "is an *F*  $\rightarrow$  is a *G*." Things are not as simple as that harmless-looking locution seems to presuppose. Stating the qualification for natural languages is a difficult matter. (Its statement in languages with canonical logical forms is straightforward: see Brandom 1994, 398.)

dard” occurrences),<sup>10</sup> those in which “is brown” occurs in conjunction with “is a cow.” Speakers who, like us, can distinguish occurrences of the former sort will have no difficulty in distinguishing, among them, those that are of the latter sort: all they need to be able to do is to recognize when one predicate occurs in conjunction with another.

Finally there is an objection to the very form of the proposal. Those who think in terms of the traditional concept of meaning would reject as counterintuitive any insertion of “dangerous, if brown” into a specification of the meaning of “cow” – which is how in their idiom the proposal I am making should be described. They would be correct to do so: we just do not use conditionals in this way when we are explaining the meaning of a word. But as I argued above, it is not clear that the inferential-role theorists who are targets of Fodor’s and Lepore’s objection are offering their theories as theories of meaning as ordinarily conceived. (Fodor’s and Lepore’s objection is that inferential roles are not compositional; their objection is not that inferential roles do not correspond to meaning *as ordinarily conceived*.) These theorists can take it as their goal to describe the inferential roles of syntactic simples so as to be able to derive from them the inferential roles of syntactic complexes, thereby making it unmysterious that creatures with finite computational capacities can master the latter. They can be indifferent to the fact that such descriptions seem odd, even implausible, when stated in terms of the traditional concept of meaning. (In a concessive spirit, however, these inferential-role theorists can point out that if we exclude the odd-seeming descriptions the remaining theory is one that does explain just those inferences that figure in our ordinary intuitions about meaning.)

I hope that the foregoing remarks satisfactorily address the “Huh?” objection. There is one final issue though.

### 3. A remaining issue

The reader will have noticed that we could just as easily have derived the data differently. We could have said that there is a SMSIC relating “a cow” and “dangerous” whose *A* set includes frames such as “Bertha is  $\alpha$  and is brown,” “The thing in the pen is  $\alpha$  and is brown,” and so on. Here the idea is that “a cow  $\rightarrow$  dangerous” inferences are proper just in case “a cow” occurs in conjunction with “brown.” The question then is this. If the “brown cow  $\rightarrow$  dangerous” inferences can be derived either from an inferential role assigned to “brown” or from one assigned to “cow,” which derivation is better?<sup>11</sup> As far as compositionality is concerned the answer, of course, is that neither is better. The situation is like that in logic, where there are infinitely many sets of axioms and rules of inference that generate the same theorems.

Whether this is objectionable depends on what one wants a theory of meaning to do. One aim, articulated by Davidson, Dummett and others, is to articulate a theory that “explicitly states something knowledge of which would suffice for interpreting utterances of speakers of the language to which it applies” (Davidson 1976, 171; see also Dummett 1975 and 1976). There is no more reason to think that that aim is met by only one theory for each language than there is to think that for each kitchen appliance there is only one user’s manual that could enable its owner to use it. In particular, there is no reason to think that someone using a theory of meaning that derived the “brown cow  $\rightarrow$

<sup>10</sup> Among the nonstandard occurrences are those in the “Gerald” sentences a few paragraphs back.

<sup>11</sup> Maybe this should be considered part of the “Huh?” objection – I just don’t know.

dangerous” inferences from an inferential role assigned to “brown” is any less able to interpret utterances in the language than someone who derived the same inferences from an inferential role assigned to “cow.” A speaker of the former sort will make precisely the same inferential moves as his counterpart of the latter sort. It is hard to see any respect in which he is nevertheless a better, or worse, interpreter than the latter.<sup>12</sup>

There is, then, at least one popular conception of the aim of a theory of meaning, with respect to which the indeterminacy we have uncovered is unproblematic. Of course there are additional aims one can have for a theory of meaning. One could require that it mirror the neural processing of language, for example. All that would mean, though, is that some theories that fare equally well with respect to the demands of compositionality fare unequally well with respect to these further *desiderata*. But then it is the further *desiderata* that are excluding some of the theories; they are not being excluded merely by the fact that two (or more) of them satisfy the demands of compositionality equally well.

Of course there is a more specific worry one might have, which is that there simply should be a fact of the matter about what the words “brown” and “cow” mean, that determines which of them is the one whose inferential role determines the propriety of “brown cow → dangerous” inferences. The reply to this has already been made: it is that “meaning,” as traditionally conceived, simply does not concern inferences other than those traditionally labeled “analytic.” Someone who thinks in terms of the traditional notion of meaning should not be concerned about the indeterminacy we are discussing, because for him it is not a fact about the *meaning* of “brown cow” that it supports inferences to “dangerous” – therefore, not a fact that calls for explanation in terms of the *meanings* of its components “brown” and “cow.”

In summary, then, the inferential-role theorist who wants to characterize inferential relations among sentences in a way that demystifies finite creatures’ grasp of them can indeed give atomistically compositional derivations of the relations in Fodor’s and Lepore’s examples, should they be demanded. But, to repeat, it is not even clear whether such a theorist is obliged to produce such derivations.

#### REFERENCES

- BLOCK, Ned 1986, “Advertisement for a semantics of psychology”, *Midwest Studies* 10, pp. 615-78.  
 BOGHOSSIAN, Paul 1996, “Analyticity reconsidered”, *Nous* 30, pp. 360-91.  
 BRANDOM, Robert 1994, *Making it Explicit*, Cambridge, MA: Harvard University Press.  
 DAVIDSON, Donald 1965, “Theories of meaning and learnable languages”, in: Yehoshua Bar-Hillel, ed., *Logic, Methodology and Philosophy of Science: Proceedings of the 1964 International Congress*, Amsterdam: North-Holland, pp. 383-94. Reprinted in *Inquiries into Truth and Interpretation*, Oxford: Clarendon Press, pp. 3-15.

<sup>12</sup>The indeterminacy in question is similarly unproblematic with respect to the aim mentioned earlier, of a theory’s explaining the possibility of a finite creature’s learning the language. A theory’s having a dual that fares equally well with respect to compositionality does not make it any less enlightening on that score.

- DAVIDSON, Donald 1976, "Reply to Foster", in: *Inquiries into Truth and Interpretation*, Oxford: Clarendon Press, pp. 171-79.
- DUMMETT, Michael 1975, "What is a theory of meaning?", in: Samuel Guttenplan, ed., *Mind and Language*, Oxford: Oxford University Press, pp. 97-138. Reprinted in *The Seas of Language*, Oxford: Clarendon Press, pp. 1-33.
- DUMMETT, Michael 1976, "What is a theory of meaning? (II)", in: Gareth Evans and John McDowell, eds., *Truth and Meaning*, Oxford: Clarendon Press, pp. 67-137. Reprinted in: *The Seas of Language*, Oxford: Clarendon Press, pp. 34-93.
- FODOR, Jerry and Ernest LEPORE 1991, "Why meaning probably isn't conceptual role", *Mind & Language* 6, pp. 329-43.
- FODOR, Jerry and Ernest LEPORE 1992, *Holism: A Shopper's Guide*, Oxford: Blackwell.
- FODOR, Jerry and Ernest LEPORE 2001, "Brandom's burdens: compositionality and inferentialism", *Philosophy and Phenomenological Research* 63, pp. 465-81.