
OUT OF CONTEXT

JERRY FODOR AND ERNIE LEPORE
RUTGERS UNIVERSITY

1. INTRODUCTION

It's been, for some time now, a pet thesis of ours that compositionality is the key constraint on theories of linguistic content. On the one hand, we're convinced by the usual arguments that the compositionality of natural languages¹ explains how L-speakers can understand any of the indefinitely many expressions that belong to L.² And, on the other hand, we claim that compositionality excludes all "pragmatist"³ accounts of content; hence, practically all of the theories of meaning that have been floated by philosophers and cognitive scientists for the last fifty years or so. A number of objections to our claim have been suggested to us, but none that we find persuasive (see, for example, the discussions of the "uniformity principle" and of "reverse compositionality" in Fodor and Lepore 2002). These objections have a common thread: they all grant that mental and linguistic content are compositional but challenge the thesis that compositionality is incompatible with semantic pragmatism. In this paper, we want to consider an objection of a fundamentally different kind, namely, that it doesn't matter whether compositionality excludes semantic pragmatism because *compositionality isn't true*; the content of an expression supervenes not on its linguistic structure⁴ alone but on *its linguistic structure together with the context of its tokening*.⁵

Here's the general idea: by stipulation, a sentence of L is compositional if and only if a (canonical) representation of its linguistic structure encodes all the information that a speaker/hearer of L requires in order to understand it.⁶ This means that, if L is compositional, then *having once assigned a linguistic representation to a sentence token, there is no more work for a hearer to do* in order to understand it. And since having knowledge of the syntax of the sentences in L and of the meanings of its lexical items is presumably *constitutive of being an L-speaker/hearer*,⁷ it follows that anyone who is a speaker/hearer of L is *thereby* guaranteed to be able to interpret an utterance of any of its sentences.⁸ The notions "speaker/hearer," "semantic interpretation," "compositionality," and "understanding" are thus inter-defined: an L-speaker is somebody who is able to understand (tokens of) L-expressions; to understand an L-expression is to grasp its semantic interpretation; an expression is compositional if and only if its semantic interpretation is determined by its linguistic structure; and a representation of the linguistic structure of an expression is "adequate" only if compositionality determines its semantic interpretation. The

upshot is that, if compositionality is assumed, *there is a definite point at which the business of understanding a token of an expression terminates*: it terminates at the assignment of whatever semantic interpretation its linguistic structure determines.

Well, the present objection is that, by these standards, English simply isn't compositional.⁹ There would seem to be lots of cases where you need to know more about a sentence token than its linguistic structure in order to interpret it:¹⁰ *you also need to know things about the context of the tokening*. For example, it's notorious that natural languages contain (perhaps ineliminably) such deictic expressions as "here," "now," and the like; and it's plausible that interpreting utterances of such expressions requires access to information about the context of their utterance. So, you need to know more to interpret an utterance of 'It's raining here' than what your knowledge of English tells you. You also need to know where the *here* in question is.

This begins to look like a dilemma. On one side, there are the usual productivity/systematicity arguments for compositionality, and there's the plausible idea that what distinguishes L-speakers as such is their ability to recognize the compositional structure of the sentences of L. But, on the other side, there's the argument that sentence tokens have their content only in their context; and, though the relevant contextual information is generally accessible to both partners in a speech exchange, still it's accessible to them qua partners in *that* speech exchange, not qua speaker/hearers of the same language. Given the impact of context on interpretation, it appears that "L-tokens are compositional" and "L-speakers ipso facto know how to interpret L-tokens" can't both be true, which is exasperating. So now what?

The literature suggests various options one might explore. We will discuss briefly a couple of them by which we are unmoved; then we'll turn to the alternative we prefer.

2. DAMAGE CONTROL

Suppose context effects show that, strictly speaking, natural languages aren't compositional; hence, that it isn't true, strictly speaking, that L-speakers as such know everything that's required to interpret L-tokens. Even so, one might say, there's a next-best possibility which, if correct, preserves the traditional story in spirit, though not to the letter. It's that the ways in which contextual variables are able to affect sentence interpretation can be demarcated a priori (that is, an exhaustive specification of the contextual variables that can affect interpretation is specified by "general linguistic theory"¹¹ (GLT)). So, it's part of what natural languages have in common 'as such'; presumably, therefore, part of the biological endowment in virtue of which any of us is able to learn to speak any of them. For example, it might be that GLT allows the interpretation of a sentence to depend on the interpretation of its constituent demonstratives and that the interpretation of demonstratives is allowed to depend on contextual information that is *not* shared by speaker/hearers as such. Still, GLT might constrain the ways in which the interpretation of demonstratives can contribute to the interpretations of their hosts. (Perhaps, for example, demonstratives

are all ipso facto singular terms.) And, likewise, GLT might constrain the kinds of contextual information that the interpretation of a demonstrative can depend on. (Perhaps demonstrated objects are required to be perceptually accessible to both members of a speech exchange.) That being so, the semantics of sentences that contain demonstratives is, as it were, *almost* compositional: their contextual interpretation depends, in ways that speaker and hearer know about in virtue of sharing a language, on contextual information that speaker and hearer know about in virtue of sharing a speech scene.

Well, if that is more or less right about demonstratives, maybe some similar treatment will work for other contextual effects on interpretation. A lot would depend, of course, on how much information the structural descriptions of sentences actually carry; and that is an empirical issue about which rational linguists may rationally disagree. Consider an utterance of ‘It’s raining’, and suppose that a good semantics for English would assign some such truth condition as *an utterance of ‘It’s raining’ is true if and only if it’s raining at the location of the utterance*. How one goes about assigning this truth condition, and how compositional the utterance turns out to be, would depend inter alia on just what the constraints the “abstract” linguistic description of the sentence is required to specify. Maybe, for example, the structure of the sentence “at LF”¹², (or wherever) is something like ‘IT’S RAINING HERE’. (So, ‘It’s raining’ and ‘It’s raining here’ turn out to have the same LF descriptions according to this proposal.) Then the only contribution of the speech scene to the truth condition of the utterance is to interpret the “underlying” demonstrative ‘HERE’. In particular, it *isn’t* required to provide the information that tokens of ‘It’s raining’ do make covert reference to the location of the speech scene. That follows just from the linguistic description of the sentence uttered since, on the one hand, LF represents ‘HERE’ as one of the sentence’s (abstract) constituents; and, on the other, the lexicon of English presumably says that ‘HERE’ means *at the location of the speech scene*. In effect, the proposal we’re considering is that all sentences are “eternal” except in respects that GLT is required to enumerate.

We will refer to this sort of theory as “Platonistic” since the basic thought is that being eternal is, as it were, the Ideal to which sentences aspire. We don’t know whether this kind of account will work in the general case (or, for that matter, even for the case of demonstratives). But we do think it’s more or less what many linguists have in mind as their model of compositionality. We have, in any event, no principled objection to pursuing a universal, a priori delimitation of the kinds of effects that context can have on interpretation, thereby rendering substantive the claim that compositionality holds except for the exceptions. On the other hand, we don’t know of anything principled in favor of the project. Let’s, therefore, consider some other options.

3. CONTEXTUAL NIHILISM

It's possible to believe, on the one hand, that context effects on interpretation preclude strict compositionality and, on the other hand, that limiting the damage by providing an a priori taxonomy of such effects is a sort of Platonist's pipe dream. The suggestion is that the effects of context on interpretation are indefinitely subtle, rich, and various; hence, that the central commitments of Platonism can't be sustained. There *isn't*, to put it in the terms we used before, "a definite point at which the business of interpreting a sentence token terminates." One just stops when the demands of communication have been met to the mutual satisfaction of both speaker and hearer. There is, in particular, no proprietary information possession of which *constitutes* the knowledge of L and which therefore distinguishes people who belong to the community of L-speakers from everybody else. Notions like "English," "English sentence," "English speaker," "learning English," and "knowing English" are thus *unprincipled* even granting the usual idealizations from idiolectic variations. Nor do the procedures for interpreting the expressions of L comprise an algorithm for the compositional analysis of utterances of its sentences. Indeed, there are no such procedures since there is no interesting sense in which sentences compose: the right model for communication is hermeneutic, not computational. We suppose that lots of (neo)-Wittgensteinian philosophers hold this sort of view; as do lots of linguists who prefer the pursuit of informal pragmatics to that of formal semantics.

So, for example, here is what we imagine Neo-Wittgenstein might say about the putative compositionality (i.e., the putative context independence) of deictic sentences like (1) and (2):¹³

(1) It's raining in New York.

(2) It's raining in Chicago.

"To be sure, if you assume that 'its raining' has the same, context invariant, meaning in (1) and (2) (viz., that it means it's raining in both), it's not all that surprising that the meanings of (1) and (2) appear to be compositional; i.e. it appears that the semantic difference between them is just the difference between a reference to Chicago and a reference to New York. But that just goes to show how much the appearance of compositionality rests on the assumption that the meanings of their lexical constituents are context invariant. Qua Neo-Wittgensteinian, I propose not to grant this assumption."

Consider (3), an example that Charles Travis has recently explored (Travis 2000). Suppose, for reductio, that (3) is compositional. (We ignore the demonstrative.) Presumably, that

(3) This ink is blue.

would require the meaning of 'blue' to be context independent.¹⁴ A natural Platonist story would be that (barring idioms like 'feeling blue' and perhaps barring intensional contexts), whenever 'blue' occurs in an English expression, it introduces a reference to the property of being blue.

But assuming that ‘blue’ always introduces a reference to being blue is sufficient to make ‘blue’ context independent only if there is a unique property of being blue. Well, is there? “The structure [of (3)] predicates being blue of some ink. That is all the structure *clearly* [sic] in the given words ‘It’s blue’. But then, in [saying] *that* [sic], one might say any of many things” (Travis, 2000, p. 197). Sometimes, it’s how the ink looks in the bottle that decides whether the ink is blue; but sometimes it’s how the ink looks on the page. And so on. The compositionality thesis, as Platonists understand it, is thus a sort of sleight of hand. That (3) is compositional depends on the context independence of ‘blue’; but the context independence of ‘blue’ depends on the assumed ontological homogeneity of being blue, of which latter neither explication nor justification is provided. So semantics and metaphysics take in one another’s wash; the myth of context-free, compositional interpretation is one with the myth of the ready-made world.

Correspondingly, according to Neo-Wittgenstein, there is no end to interpretation:

...take any statement [S] and ascribe to it any set of representational features [F] you like...[If] we then find a statement with F which still differs from S in when it would be true...we assign S another feature F*...But in the envisioned situation, no matter how we start, or how we continue this process, there is no way of bringing to a halt the sequence of statements which, sharing more and more representational structure with S, nonetheless differ in content¹⁵ (ibid., p. 36).

If that’s how language works, then content is inherently context dependent and interpretation is hermeneutic all the way down. Or rather: interpretation is hermeneutic all the way down *because* content is inherently context dependent. The situatedness of meaning is one with the ineliminability of interpretation.

We’re not, of course, endorsing this view; so far, indeed, we’re not endorsing anything. Though neo-Plato has lots of problems, neo-Wittgenstein does too. As we mentioned above, it appears there are aspects of language (/thought) that compositionality is essential to explaining.¹⁶ Moreover, the reduction of semantic interpretation to hermeneutics isn’t intuitively very plausible. Figuring out what somebody said really does seem quite different from figuring out what he might have meant by saying it. The former is generally fast, fluid, and automatic; the latter is often a matter for endless reconsideration. If someone utters ‘It’s raining’ in the usual conditions, he thereby says that it’s raining; if you don’t believe us, ask your local English speaker. But the paradigm of hermeneutics is text interpretation, which goes on forever (as do, in consequence, departments of literature). It’s pretty clear what ‘It’s raining’ means, but God only knows what ‘Hamlet’ means (or even ‘As You Like It’). The jury was still out as of this writing.

So much for two standard ways of thinking about how context and linguistic structure might interact when a language is used for communication. We turn now to what we care about most. We’re going

to argue that there's a reason—indeed, a glaring one—why neither Platonism nor Wittgensteinism could conceivably be true.

4. THE ASYMMETRY ARGUMENT

Both the views we've summarized take for granted that the objects to which semantic interpretations are assigned are (not linguistic utterances *per se*, but) linguistic utterances together with their contexts. This is obvious in the case of neo-Wittgenstein; that utterances have interpretations only in context is the very burden of his plaint. But it's also true for neo-Plato since it is untendentious even in his camp that the content of utterances is sensitive in various ways to contextual determinants, deicis and ambiguity resolution being the clearest cases. As previously remarked, the Platonist project is not to rid interpretation of context dependence; it's only to delimit *a priori* which contextual parameters can be germane.

So neo-Plato and neo-Wittgenstein are both committed to the view that what really gets interpreted when language is used are pairs of sentence-tokens and contexts. Moreover, they both hold this not just as epistemology but also as metaphysics. This point is crucial. It's an epistemic truism that, if how a certain utterance is to be understood depends on some fact about its context, then a hearer can't understand the utterance unless he is apprised of that fact. But it doesn't follow that such epistemologically essential facts are metaphysically constitutive of an utterance meaning what it does. And, as we read them, it's precisely this metaphysical claim that neo-Plato and neo-Wittgenstein agree about. They both think that the interpretation of an utterance supervenes on (*inter alia*) the context of its utterance. We take it that supervenience is a metaphysical relation.

So then to summarize: neo-Wittgenstein and neo-Plato both think not just that information about context can bear on the hearer's interpretation of the speaker's utterance, but that some (or many or all) of the properties of the context are constitutive of the utterance's having the interpretation that it does. We take it that the second claim just doesn't follow from the first; anyhow, it doesn't follow lacking a lot more argument. From the fact that you can (reliably and with warrant) infer that it's raining from puddles in the street, it just doesn't follow that whether it's raining supervenes on whether there are puddles. Epistemology is one thing, metaphysics is another; nothing good can come of conflating them, either by accident or as a matter of principle.

Here's the view we favor (for want of a better term, we'll call it "Cartesian"): nothing about the context of an utterance is a metaphysical determinant of its content. The only metaphysical determinants of utterance content are (i) the linguistic structure of the utterance (the syntax and lexical inventory of the expression type that it's a token of), and (ii) the communicative intentions of the speaker.¹⁷ Nothing else. Ever. What exactly this means, and why we think it must be true, is the rest of what follows.

Let's start with ambiguity resolution, since this seems, *prima facie*, to be a clear case where facts in "the background"¹⁸ can determine

what an utterance means. But we'll argue, on one hand, that this can't be metaphysical determination; and, on the other, that what we have to say about the resolution of ambiguity holds, *mutatis mutandis*, for the other putative instances where semantic interpretation is background-sensitive.

Groucho said, as everybody knows, "I shot an elephant in my pajamas." This sets up the infamous joke: "How an elephant got into my pajamas I can't imagine. [Laughter]." What, exactly, happened here? We take the following to be untendentious as far as it goes: the conventions of English are in force, and they entail that there are two ways to read the set-up sentence. Either it expresses the thought (I, in my pajamas, shot an elephant) or it expresses the thought (I) (shot (an elephant in my pajamas)). Giving the communication context (including the operative background of shared beliefs about elephants, pajamas, and so forth), it's natural to interpret what Groucho said according to the first parsing rather than the second. But what he says next shows he meant that the pajamas were on the elephant, and that pulls the rug.

So, background considerations can determine the truth conditions of an (otherwise) ambiguous utterance, and the issue is whether they do so by providing metaphysically necessary conditions for the utterance to have the interpretation that it does. We think not. We think that what makes one or other disambiguation the right one is only that part of the context that consists of what Groucho had in mind. Accordingly, the rest of the background has only epistemic relevance to the interpretation of his utterance: it provides information that (e.g.) Harpo can use to figure out how Groucho intended the utterance to be understood. (In the present case, the background is misleading as to what Groucho intended; that's the joke, don't you see.)

There are three, closely connected, arguments for this way of understanding the case. First, if the facts about the background were constitutive of the correct disambiguation, it would presumably follow that there are contexts in which, in point of metaphysical necessity, Groucho couldn't make his joke; that is, contexts in which the right interpretation of the set-up sentence couldn't be that the pajamas were on the elephant. But we take it that there are no such contexts; the metaphysics of disambiguation places no constraints at all on Groucho's wit. This is immediately evident if, as we suppose, the contribution of the background is merely evidential. No doubt there could be situations in which the context virtually guarantees that the speaker was saying this rather than that. But "virtually guarantees" is an epistemic concept; one that metaphysics knows not of. What is, from an epistemic point of view, virtually guaranteed, may nonetheless not be the case.

Here's a second way to make the point. The possibility we're considering is that only utterances-in-contexts have meanings, hence that Harpo can know which way to parse Groucho's utterance only if he (Harpo) is apprised of the relevant contextual facts. But likewise for Groucho, or so it would seem. If only contextual information can resolve the ambiguity of the utterance, then someone not apprised of that

information can't parse it, and this must apply to Groucho *inter alia*. Barring access to relevant contextual information, Groucho himself can't know whether what he said was true if and only if (he in his pajamas) (shot an elephant) or whether it's true if and only if (he) (shot (an elephant in his pajamas)). But that's preposterous on the face of it. It's mad even by the Marx brothers standards because it leaves room for a scenario in which not just Harpo, but also Groucho, hunts around in the context to find the facts that determine what it was that Groucho said about the elephant. Patently, there can be no such scenario.

To put it still another way: it's pretty plausible that, in some sense or other, Groucho has privileged access to his communicative intentions. But, then, his communicative intentions surely can't be constituted by (or even metaphysically dependent on) what is contingently going on in the context; if they were, then Groucho's privileged access to his communicative intentions would imply correspondingly privileged access to the environmental contingencies. But nobody—not even Groucho—has privileged access to environmental contingencies. Knowledge of environmental contingencies is always *a posteriori*, hence not privileged; as, indeed, Hume made abundantly clear.

These are, we suppose, three variations on the same theme. The upshot, in each case, is that there is an inherent asymmetry between the epistemological situations of the speaker and the hearer with respect to the role of contextual information in the disambiguation of an utterance: the hearer can use such information but the speaker can't. That's part and parcel of the fact that the speaker, but not the hearer, has immediate (privileged, non-inferential) knowledge as to which disambiguation is the right one. It's clear how this could all be so if disambiguation supervenes on the speaker's intentions since, for better or worse, one's intentions just *are* the sort of things to which one's access is typically privileged.¹⁹ But how could a speaker (or anybody else) have immediate, non-inferential access to, as it might be, the fact that elephants don't wear pajamas?²⁰

If, in short, disambiguations supervene on facts about the background, then it would seem that one's disambiguating interpretations would always be inferences from one's grasp of such facts. But Groucho's access to the truth conditions of his utterance isn't inferred from what he knows about the background; indeed, it isn't inferred at all. That's all as it should be if disambiguation supervenes on the intentions of speakers, but we can't see how to make sense of it on any other assumption.

So, then, unless disambiguation is atypical of interpretation at large,²¹ the moral would seem to be that, since the speaker's access to the interpretation of his utterance is epistemically privileged, nothing about the background of an utterance is metaphysically constitutive of its interpretation. The function of background knowledge in interpretation is (only) to provide premises for the hearer's inferences about the speaker's intentions.²² If the resolution of ambiguity is typical of interpretation at large, then what has content is not speech-in-a-context but speech as its speaker intends it.

This line of thought actually isn't novel. So, Evans writes:

...when an audience or a theorist confronts an utterance of an ambiguous sentence, the only sensible direction in which he may look for information enabling him to disambiguate the utterance is towards facts which bear on the speaker's intentions. This might lead one to think that it is at least a *necessary* [sic] condition for saying that p that the subject have the intention to express the thought that p; and as this would ordinarily be understood, it would require the subject to have, or at least to be capable of having, the thought that p (Evans, 1982, p. 68).

One might indeed be lead to think that. What, then, does Evans think is wrong with drawing the inference? Here's how he continues:

This principle [saying p requires having, or being able to have, the thought that p] would thus seem to legitimize delving into the half-baked ideas and misconceptions people have associated with at least the more specialized words of the language, in order to decide what a speaker is saying when he utters the words [4].

4. This ship is veering to port'. (*ibid*; renumbered).

We're not entirely clear what the argument here is supposed to be; but perhaps the idea is this: if what you said depends on what you were thinking when you said it, then what you said when you said 'port' depends (*inter alia*) on whether you were thinking about the ship as veering to the left or as veering in the direction of the intended destination.²³ But it's by no means certain just what it is that you must have in mind in order to mean 'port' one way or the other. For example, (as a matter of fact) to turn to port (= left) you must turn in the direction that is to the left when you are facing forward. So, then, how you should interpret an utterance of 'port' would seem to depend on what you assume the speaker knows about (how much you assume he knows about) this piece of nautical jargon. Correspondingly, the steersman to whom the skipper says "turn to port" might reasonably complain not just that it's unclear what he's been told to do, but that he has no way of finding out short of an inquiry into the skipper's psycholinguistic situation. This seems counterintuitive (it wouldn't get you off the hook with any skipper that we've crewed for). In any case, it's a view one might well wish to avoid.

Evans has a story about how to avoid it: disambiguation is special; it depends on the speaker's intentions, but the rest of interpretation doesn't. "Once it is clear which linguistic counter [the speaker] is putting forward [viz., which reading of the ambiguity he intends] the content of what he says is determined by the significance which that counter has in the game, and not by whatever half-baked and ill-informed conception he may have of its meaning" (Evans, 1982, p. 69). So, once it's fixed that the relevant 'port' is the one that's opposite of 'starboard' (rather than the one that's in 'home port'), what the skipper ordered the helmsman to do was to turn to the left of the boat. This is so whether or not the skipper knows about the "facing forward" rule (indeed, it's true even if the helmsman knows that the skipper doesn't know about the "facing forward" rule. Presumably that's because the helmsman

and the skipper must both defer to expert users of the term.) A methodological intuition that's calling the shots here is that "...in general, the notion of understanding [sic] a reference of a certain type is a more fundamental notion than the notion of making a reference of that type" (Evans, 1982, p. 171). And the ontological assumption that underlies the methodology is that a language is, inherently, a social object "Once one's interest is in the phenomenon of language itself, one must be concerned with the way in which it functions as a means of communication among speakers of a community" (Evans, 1982, p. 67).

Well, maybe so; but this line of argument does rather beg the question against the Cartesian view that the function of language is to express thoughts, and it's what thought an utterance is intended to express that determines what interpretation of the utterance is right. (Also, for what it's worth, our intuitions about Evans's example are unclear. In all the boats we've been on, the crew is well advised to defer to the skipper, and the skipper defers only to God.) In any case, we think that Evans has gotten hold of a red herring. On our view, there are actually two quite different issues that arise in the context of interpretation. One is: what the speaker intended to say; the other is: how he ought to have said it. We think that speakers are privileged with respect to the first sort of question, but not with respect to the second. Solecism is always possible; that applies not just with respect to a context dependent expression but to any expression at all. So, Evans was right to say that the social character of language "immediately opens up the possibility of gap between what a speaker means to say by uttering certain words—what thought he wishes to express—on the one hand, and what he strictly and literally says, according to the conventional meanings of the words he utters, on the other" (Evans, 1982, p. 67). But though that's true enough, it shows a lot less than Evans seems to have supposed. In particular, it *doesn't* show that (once ambiguities are resolved) the right interpretation of an utterance depends on the linguistic conventions that determine what the speaker "strictly and literally says."²⁴ The moral is just that, in the hearer's pursuit of an interpretation, the possibility of a solecism on the speaker's part has got to be kept in mind.²⁵

This is all unsurprising from the Cartesian point of view. The speaker's beliefs about his communicative intentions are privileged; but, of course, his beliefs about English aren't. Since English is, in some sense, a *public* institution, whether one has conformed to the linguistic norms is, in principle, no more at one's discretion than whether one has conformed to the whims of the IRS. Even *ignorance* of the rules doesn't, in either case, count as a defense. Indeed, Cartesians think that (in the usual case) it's part of the speaker's *intention* that his utterance should be interpreted in a way that conforms to the grammar (syntax, semantics, whatever) of the language that he and his hearer share. *That* intention is privileged; you can't be mistaken about whether it was English or German that you were speaking when you uttered "Emedokles leaped." But your conviction that you actually did say what you intended to is always corrigible in principle.

We think that, solecism to one side, the speaker is always privileged as to what he intended to say, just as he is always privileged as to how he intends his ambiguities of utterance to be resolved. In this respect, demonstration (for example) works just like disambiguation. Contextual features can't be constitutive determinants of what the speaker demonstrates; for, if it were, there could be a scenario in which the speaker *finds out* what it is that he has demonstrated by finding out about these features of the context. ("I thought that in saying 'that ear' I had demonstrated my left ear. But now I see that it was my right ear after all.") But that seems mad; I can, in fact, demonstrate my left ear in any context in which I have one. According to Cartesians, this is a consequence of my freedom to *think about* my left ear in any context in which I choose to do so. It is of course true not withstanding, that though the object of my demonstration is in fact my left ear, there are contexts in which my audience is entirely justified in supposing otherwise. Suppose my right ear is patently on fire and my left ear patently is not. Still, I can, if I choose, utter 'That ear is on fire', thereby demonstrating my *left* ear (and thereby saying something patently untrue). What I can't reasonably do is complain about my demonstration having been misconstrued.

To be sure, all this turns on appeals to mere intuition and we admit that other English speakers may find ours to be idiosyncratic. Here's a passage (slightly altered) from David Kaplan:

Suppose that without turning and looking [Groucho points] to the place on my wall which has long been occupied by a picture of Rudolph Carnap and [says]: Dthat... is a picture of one of the greatest philosophers of the twentieth century. But unbeknownst to [Groucho] someone has replaced my picture of Carnap with one of Spiro Agnew. I think it would simply be wrong to argue an 'ambiguity' in the demonstration, so great that it can be bent to [Groucho's] intended *demonstratum*. [Rather, Groucho has] said of a picture of Spiro Agnew that it pictures one of the greatest philosophers of the twentieth century. And [his] speech and demonstrations suggest no other natural interpretation to the linguistically competent public observer (Kaplan, 1978, p. 355).

We aren't deeply moved by this.²⁶ It seems Kaplan is offering the point about what interpretation is natural for the linguistically competent observer to assign as an argument that Groucho has (inadvertently) demonstrated Agnew. If so, that strikes us as begging the question. Granting that, in the situation Kaplan describes, no other natural interpretation is available to the linguistically competent public observer as such. But what's at issue is precisely whether, in such cases, the correct interpretation of an utterance *is* available to linguistically competent public observers as such. Perhaps, it's available only to the linguistically competent speaker who is clued in to what Groucho thinks about Agnew, Carnap, and about the philosophical enterprise at large. To suppose that what linguistically competent speaker/hearers know as such determines the correct interpretation of the speaker's utterance is to take sides against the Cartesian suggestion that only what the speaker intends does. So, why prefer Kaplan's analysis to one that says

that, in the situation imagined, the linguistically competent public observer is bound to mistake the object of demonstration? That this mistaken interpretation would be fully justified doesn't, we suppose, show that it would be *true*. Granting that Groucho could have been wrong about *what he was pointing at* when he said "That is a picture of one of the greatest philosophers of the twentieth century," could he likewise have been wrong about *whom he was referring to*? Or (equally bizarre to our ears) could Groucho have been referring to one guy and his utterance to another?²⁷

One last thought along these lines before we sum up. We've often been told that, if we adopt the Cartesian view, we're going to have a hopeless problem with indexicals. Suppose, for example, that John, who (falsely) believes himself to be Napoleon, says "I'm sick," intending thereby to say of Napoleon, that he is sick. Still, the truth maker for his utterance implicates John's health, not Napoleon's. Doesn't that show that, with indexicals even if not with demonstratives, the context of utterance overrides the speaker's communicative intentions?

Well, that's one way to analyze the case; but here's another: 'I' is a singular term, just like 'this'. The rule for both is that, if a singular term is used to report a demonstration, the reporter must demonstrate the same individual that, the reported speaker did. So, John says 'I'm sick' thereby referring to John. I can report what John said with "John said that he's sick," so long as my 'he' also refers to John. The difference between indexicals and singular terms like names, demonstratives, etc., isn't that indexicals let context override communicative intentions; rather, it's that indexicals don't survive disquotation when they are reported. You say "This is pink." I can report what happened by saying "He said this is pink" so long as my 'this' demonstrates whatever yours did. But indexical demonstrations don't survive disquotation, so I can't report your utterance of "I'm sick" by saying "He said that I'm sick."

Q: 'But if 'this' and 'I' are both singular terms, how come reports can disquote the one but not the other?'

A. Because 'I', unlike 'this', is a dedicated term; 'this' picks out whatever the speaker demonstrates, 'I' always picks out the speaker.

Q. Well, but when John said "I'm sick" John's intention was to communicate that Napoleon is sick; but his 'I' went back to John all the same. Isn't that a clear case where a communicative intention gets over-ridden by a context of utterance?

A. No. You are assuming what we don't intend to grant: that there are no indexicals in thought. What John was intending to communicate, according to us, was that he's sick. This intention is exactly preserved by reporting that John said that he's sick, and that's so whether or not John thinks he's Napoleon. One mustn't overlook that "intends to say that" is an opaque context; one in which substitutivity of coreferentials fails. In particular, it's a context in which you can't substitute an indexical for a name *modulo* what one intends to say; not even if you take them to corefer.²⁸ John's "I'm sick" expresses his thought *that's he's sick* whether or not John thinks he is Napoleon.

The moral of all this is pretty unamazing: in a case when a speaker's gesture (verbal, ostensive, or whatever) may be inappropriate to what he intended to communicate, the interpreter must decide whether to go by what was actually uttered or by what, in the circumstances, the speaker is likely to have meant to say by uttering it?²⁹ And, since there's no principle that applies in the general case, it's not surprising that intuitions about which is the right decision are often labile and very sensitive to the particulars of examples. Still, according to us, if the speaker's intention is given, so too is the correct interpretation of his utterance; that's so even though the context is irresistibly misleading.³⁰ Epistemology is one thing; metaphysics is quite a different thing. Or perhaps we've said that.

5. BUT AFTER ALL, SO WHAT?

Perhaps you are feeling a bit disappointed? Well, we sympathize. We promised to make a case that no facts about context are constitutive of linguistic content, and we now claim to have done so. But we haven't denied the effect of contextual variables on the content of *thought*, nor do we propose to. On the contrary, we suppose an utterance of 'That's on fire' inherits its truth conditions from the thought that it expresses; and a thought *that's on fire* is true or false depending on whether what it demonstrates is on fire. We are thus still in want of a metaphysical story about how context can determine content. All that's altered is the locus of the effect.

On the other hand, we think the difference of locus matters a lot. We'll close on that note.

The usual ways that philosophers think about the problem of context effects in language assumes, practically invariably, that it's part of the problem of interpretation. (Thus Evans's dictum quoted above: "...in general, the notion of understanding a reference of a certain type is a more fundamental notion than the notion of making a reference of that type"). But this approach to context effects ignores the speaker's epistemic privilege with respect to the objects of his demonstrations, the resolution of his ambiguities, and the like. Clearly, that needs to be fixed. As far as we can see, fixing it requires enforcing the distinction between interpretation_v, which is something that speaker/hearers *do* in the course of a communication exchange, and interpretations_N, which are things that symbols *have* (indeed, things that they have essentially). Prevailing philosophical opinion is that what *gets interpreted*, and what *has an interpretation* are both expressions in a natural (a fortiori, public) language. Indeed, it's the received view that interpretations_N are ontologically parasitic upon the practice of interpretation_v. That being so, the metaphysics of content and the epistemology of content assignment can't dissociate even in principle. So the story goes; so it's gone for years.

But (according to us) the asymmetry argument shows that can't be right. Rather, what has content in the first instance is the propositional attitudes of "intentional systems"; most notably, for present purposes, the communicative intentions of speaker/hearers. By contrast,

interpreting is the process whereby hearers recover the content of communicative intentions from the noises that speakers make when they try to express them. What *get* interpretations_V are utterances; what *have* interpretations_N are states of mind. Some metaphysical story about the content of communicative intentions must thus be *prior* to any story about the epistemology of linguistic communication.

One last point along these lines. To treat contextualization as an aspect of interpretation (i.e., to think about it, “from the hearer’s point of view” rather than the speaker’s) is perfectly reasonable if what you’re primarily worried about is the epistemology of content ascription. And it’s perfectly reasonable to be primarily worried about the epistemology of content ascription if the refutation of skepticism is high among your philosophical priorities. *Prima facie* anyhow, the epistemological situation of the ascriber is strikingly different from that of the ascribee; and it’s generally supposed that it’s only about the former that skeptical issues arise (presumably because paradigm second person mental ascriptions are *inferential* but paradigm first person mental ascriptions aren’t).³¹ It seems that there really is an asymmetry built into the epistemology of the mental, and the view from the first person’s position really is quite different from the position of the second person. We have no quarrel with this understanding of the epistemological landscape. Indeed, we think something of the sort has to be true if we’re to suppose that mental state ascriptions have truth values at all since we think that first-person mental states are the truth makers of second-person mental ascriptions.

But if it’s pretty clear that the epistemology of the mental is asymmetrical, it’s even clearer that the metaphysics (of the mental or of anything else) couldn’t be. Metaphysics is about relations like supervenience, determination, and the like, and *they aren’t perspectival*. Maybe, for example, mental states supervene on brain states; we don’t know, but we wouldn’t be entirely surprised. What’s out of the question, however, is that mental states supervene on brain states *from your point of view but not from mine*. Supervenience isn’t that kind of thing; and neither are identity, duality, epiphenomenality, eliminability, or any other of the usual candidates for connecting the mind to the brain.

Epistemic states are perspectival more often than not, but metaphysical states never are. That’s why you don’t have a chance of deriving your metaphysics from your epistemology, however hard you try.³²

Endnotes

1. In previous papers, we’ve usually assumed that the implications of compositionality are the same for linguistic and for mental representation. However, in the present discussion it matters a lot which of these is under discussion. (Why that’s so will presently emerge.) Our strategy will be to start with language and work our way in.
2. We take for granted the usual idealization from “performance” variables.
3. In our usage, a pragmatist theory of content is one that construes it in such *epistemic* terms as for example, ‘knowing how’ or ‘believing that’. For discussion, see Fodor (2004).

-
4. By stipulation, a representation of the “linguistic structure” of an expression specifies its syntax together with the semantic interpretations of its lexical primitives. So, for example, if the semantic interpretation of a sentence is something like an assignment of its truth conditions, then a specification of the linguistic structure of the sentence ‘John loves Mary’ might say that its syntax is $((\text{John})_N ((\text{loves})_V (\text{Mary})_N)_{VP})_S$ and that ‘John’ denotes John, ‘Mary’ denotes Mary and ‘loves’ denotes the relation x loves y . There are, of course, lots of other ways understanding such notions as semantic interpretation and linguistic structure; to say that philosophers disagree about which is the right one considerably understates the case. For purposes of exposition, we will generally adopt a truth theoretic way of talking about content, but we intend the polemics not to depend on our so doing.
 5. Throughout the discussion we construe the notion of context very liberally; it’s to include not just the local physical setup, but the intentional environment as well; that is, the speaker and hearer’s beliefs, desires (and so forth) insofar as they may affect the communication process. We thus mean to include the relevant aspects of “background” in something like Searle’s use of that term (Searle 1980).
 6. Notice that it’s plausible that sentences have linguistic structure at many levels of representation, which is to say that they may contain more lexical and/or syntactic material than they display on their surface. So, for example, it’s plausible that there’s a level of linguistic representation at which the syntax of ‘John expects to leave’ is something like JOHN EXPECTS (JOHN TO LEAVE). In such cases, the conditions for the compositionality of a sentence are satisfied only by their abstract representations (assuming that they are satisfied at all). We’ll say that compositionality requires that the meaning of an expression be determined by its lexical inventory together with its “complete structural description.” We mean to leave open what properties of an expression its complete structural description specifies.
 7. For present purposes, we don’t care whether, or to what extent, “what the speaker/hearer knows about this language” is known explicitly.
 8. *Modulo* the usual idealizations about performance and competence (which we shall take for granted throughout what follows.)
 9. Except insofar as it may contain “eternal” sentences (Quine’s term for sentences whose truth conditions are entirely independent of the context of their utterance; ‘ $2+2=4$ ’ and the like are plausible candidates).
 10. Note the systematic ambiguity between “interpretation_N,” which is what the linguistic structure of an expression determines, and “interpretation_V,” which is what an L-speaker does in virtue of which he understands an expression-token in L. We’ll return to this presently.
 11. That is, it’s part of what natural languages have in common “as such”; hence, presumably, part of the biological endowment that enables any of us to speak any of them as a native language.
 12. That is, at the (putative) level of logical form.
 13. Arguably, both make covert reference to the *time* of utterance. But let’s just ignore that.
 14. Notice that this is *prima facie* compatible with the thesis that the meaning of ‘blue’ is syncategorematic; for example, that ‘blue X’ always means something
-

like 'blue for an X'. Since we aren't proposing to defend either neo-Plato or neo-Wittgenstein, this concession is harmless for our purposes.

15. Travis doesn't argue—or even claim—that the “envisioned situation” actually obtains. But the context makes it pretty clear that he thinks that quite possibly it may.
16. For example, the validity of syllogistic reasoning appears to require it. *Prima facie*, the inference from ' $P \& Q$ ' to ' Q ' is invalid unless ' Q ' in the premise means the same as ' Q ' in the conclusion.
17. But doesn't the speaker's communicative intention also count as part of the context? If you like, but the issue is merely verbal. We're content to say that the contextual determinants of the content of an utterance include *only* the communicative intentions of the speaker.
18. We borrow Searle's term in order to emphasize that the context relevant to interpretation often includes the hearer's beliefs about the speaker's beliefs about the world (including, Grice-wise, the hearer's beliefs about the speaker's beliefs about the hearer's beliefs). Also, it can be what's in the *linguistic* environment (e.g., in the structure of the discourse in which the utterance is embedded).
19. Or at least one's conscious intentions are. We assume that speakers are normally aware of the intended interpretation of their utterances; Freudian slips aren't the general case.
20. Contrast a species of asymmetry argument that we aren't inclined to sign on for: “I have non-inferential access to my pains but not to my brain states; *ergo*, my pains aren't brain states.” What's wrong with that is that maybe the premise is false; maybe I do have non-inferential access to my brain states under such descriptions as 'my pain' (see, e.g., Loar 1990). But (to repeat the text) there isn't the corresponding possibility in the elephants-don't-wear-pajamas case; nobody has privileged access to those sort of facts under any description.
21. In fact, there is a good reason for supposing that the treatment of disambiguation must be homogeneous with the treatment of other kinds of context sensitivity (as, for example, indexicality). Namely, that there are cases where disambiguation itself turns on the assignment of an interpretation to a contextually sensitive item. Consider 'John said he left three days ago'. This is scope ambiguous between (roughly) *he said (he left three days ago)* and *three days ago (he said he left)*. Here disambiguation turns on the interpretation of the indexical element 'ago'; viz., on whether the three days are counted from John's 'now' or from the speaker's. In such cases, one can't coherently claim *both* that the speaker is authoritative with respect to the disambiguation of his utterance and that he is not authoritative with respect to the interpretation of the indexical constituents.
22. Notice that one can't get into the Groucho/Harpo situation when one is talking to oneself. (One says to oneself: “Flying planes can be dangerous” and wonders whether the thought was that it can be dangerous to fly planes or that planes can be dangerous when they fly. Not, we suppose, a possible scenario).
23. We don't, as a matter of fact, think that one would say either this way. As we understand the nautical dialect, it's winds, not ships, that veer. Ships *turn* (like buses). But perhaps that's not crucial.
24. This is one of the (many) respects in which the analogy between languages and games is unhelpful. There's sometimes a point, in chess, to holding a player to

the move that he indicated. But the point of interpretation isn't to hold the speaker to the letter of his utterance; it's to figure out what thought he intended his utterance is to convey.

25. Suppose someone says 'It's raining here' believing, wrongly, that he is in Pittsburgh. Then the right interpretation for the hearer to impose on the utterance is *it's raining in Pittsburgh*, not *it's raining here*. To insist on the "literal" meaning would be pointless and pedantic, not to say uncharitable. There is, in this respect, no principled difference between indexicals and other sorts of terms; that solecism is always possible is something the hearer must always bear in mind.
26. It complicates the example unnecessarily that what Groucho is (or isn't) demonstrating is supposed to be a *picture* of Agnew; since on nobody's view is it the picture that Groucho thinks was a philosopher. We'll ignore this in what follows.
27. Actually, Kaplan has changed his mind about this, and his revised view strikes us as exactly right. On p. 582 of "Afterthoughts" he says: "In 'Demonstratives' I took the demonstration, typically a (visual) presentation of a local object discriminated by a pointing, to be criterial for determining the referent of a demonstrative. While recognizing the teleological character of most pointing—it is typically directed by the speaker's intention to point at a perceived individual on whom he has focused—I claimed that the demonstration rather than the directing intention determined the referent... I am now inclined to regard the directing intention, at least in the case of perceptual demonstratives, as criterial, and to regard the demonstration as a mere externalization of this inner intention. The externalization is an aid to communication, like speaking more slowly and loudly, but is of no semantic significance."
28. It's independently fairly clear that exchanging a name for a pronoun (more precisely, for a trace) in a report can affect truth value even when the two actually are coreferential. Suppose Churchill says "Only I can remember giving the speech about blood, sweat, toil and tears." He thereby says something true since only he did give the speech. But I can't report him as having said that only Churchill can remember *his* (viz., Churchill's) giving the speech; in fact, millions of people remember his giving it. (It's true that, if what he'd said was "Only I remember myself giving the speech...", I could report him as having said "Only he remembers himself giving the speech." But this doesn't help since there's a (subtle) distinction between 'I remember giving it' and 'I remember myself giving it'; it's the former that makes trouble; the latter doesn't.) This is the only case we know of where it is demonstrably impossible to report what someone else thinks. For further discussion, see Fodor (1975).
29. To put it in our previous terms, solecism is the case where the speaker's intention that his utterance be taken to mean such and such is incompatible with his intention that it be interpreted according to the norms of the language that he and the hearer share. This makes a problem for the hearer who wishes, in general, to do the best for the speaker's intentions that he can. There isn't, we think, any general principle that decides; sometimes it's the spirit that matters, sometimes it's the word.
30. In passing: this suggests a rather different treatment of the "arthritis" cases (see Burge, 1979) than they have generally received. Roughly, we don't think they show that deference determines the content of belief. Rather, they arise because one wishes to report what is said (or, *mutatis mutandis*, what is thought) as

unmisleadingly as possible; and, in these cases, there are conflicting considerations about how to do so. If, according to the experts, arthritis is a disease of the joints, then that's what 'arthritis' means in English. But, of course, Jones doesn't know what the experts say about arthritis, and his view is precisely that something's gone wrong with his thigh. In such a case it would be misleading to say, without caveat, either that Jones thinks he has arthritis in his thigh or that he doesn't. That shows that it's sometimes hard to decide what's best to say about what someone says or thinks; it doesn't show that the content of an utterance or of a thought can depend on patterns of deference.

31. Inferential accounts of mental self-ascriptions have in fact been offered from time to time (see e.g., the "James-Lang" theory of emotion); but largely, we suspect, pour epater les philosophes. "I don't know what I believe until I hear what I say" is, maybe, mildly amusing; but, in the vast majority of cases, it simply isn't true.
32. Strawson once suggested (if we read him right) that there are second-person criteria for mental states, but no first-person criteria. This wasn't, however, the thesis that the (metaphysically) sufficient conditions for being in pain are perspectival; rather, it was the thesis that 'pain' and the like are ambiguous.

Bibliography

- Burge, T. "Individualism and the Mental," in French, Uehling & Wettstein. *Midwest Studies in Philosophy IV: Studies in Metaphysics*. Minneapolis: University of Minnesota Press, 1979.
- Chomsky, Noam (1992). "Language and Interpretation: Philosophical Reflections and Empirical Inquiry." In J. Earman (ed.) *Inference, Explanation and Other Frustrations: Essays in the Philosophy of Science*. Berkeley: University of California Press, 1992, 99-128. Reprinted in Chomsky 2000.
- _____. (1993). *Language and Thought*. London: Moyer Bell.
- _____. (1995). "Language and Nature." *Mind* 104: 1-61. Reprinted in Chomsky 2000.
- _____. (2000). *New Horizons in the Study of Language and Mind*. Cambridge, UK: Cambridge University Press.
- Evans, Gareth, 1982. *The Varieties of Reference*. Oxford. Oxford University Press.
- Fodor, J. *The Language of Thought*, 1975.
- Fodor, J. "Having Concepts," forthcoming 2004, *Mind and Language*.
- Fodor, J and E. Lepore. *Holism: A Shopper's Guide*. Basil Blackwell, 1992.
- _____. *The Compositionality Papers*, Oxford University Press, 2002.
- Kaplan, David, "Dthat," in Cole, Pete (ed.) 1978. *Syntax and Semantics*, vol. 9, New York: Academic Press, 21-53.
- _____. "Afterthoughts," in *Themes from Kaplan*. Oxford: Oxford University Press, J. Almog, J. Perry, and H. Wettstein 1989, 567-614.
- Loar, Brian, "Phenomenal States," *Philosophical Perspectives*, ed., J. Tomberlin, 1990, 81-108.
- Searle, John. "The Background of Meaning," in *Speech Act Theory and Pragmatics*, eds., Searle, J., F. Kiefer and M. Bierwisch (Reidel: Dordrecht), 1980, 221-32.
- Travis, Charles, *Unshadowed Thought*. Cambridge: Harvard University Press, 2000.
-