

Context and Compositionality

Jerry Fodor and Ernie Lepore

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 languages explains how L-speakers can understand any of the indefinitely many expressions that belong to L. And, on the other hand, we claim that compositionality excludes all 'pragmatist' 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/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 content pragmatism. In this paper, we want to consider an objection of a fundamentally different kind: namely, that it doesn't matter whether compositionality excludes content pragmatism because *compositionality isn't true*; the content of an expression supervenes not on its linguistic structure alone but *on its linguistic structure together with the context of its tokening*.

Here's the general idea: by stipulation, a sentence of L is compositional iff a (canonical) representation of its linguistic structure encodes all the information that a speaker/hearer of L requires in order to understand it. 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* 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* 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. 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 iff its semantic interpretation is determined by its linguistic structure; and a representation of the linguistic structure of an expression is 'adequate' only if it compositionally determines its semantic interpretation. The upshot is that, if compositionality is assumed, *there is a definite point at which the business of understanding 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. 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: *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 rest; and it's plausible that interpreting utterances of such expressions requires access to information about the context of the utterance.

So, if interpretations are assignments of truth conditions, then 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 precisely 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 suggest the alternative that we favor.

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' (GLT)). So, 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 speakers and hearers all know about in virtue of their sharing a language, contextual information that speaker and hearer both usually know about in virtue of their sharing a speech scene.

Well, if that is more or less right about demonstratives, maybe some similar treatment will work for other parameters of 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 iff 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 assumptions about the 'abstract' linguistic description of the sentence type that the utterance tokens. Maybe, for example, the structure of the sentence

‘at LF’, (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’ make covert reference to the location of the speech scene. That follows from just the linguistic description of the sentence uttered since, on the one hand, LF represents ‘HERE’ as one of its 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 their linguistic structural descriptions are required to specify.

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 strict, *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, 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 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 and 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 deixis-free sentences like (1) and (2):

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 context-invariant *difference* between the meanings of (1) and (2) is just the difference between a reference to Chicago and a reference to New York. However, that consideration doesn’t at all show that (1) and (2) are compositional; it only illustrates how much the compositionality of sentences rests on the assumption that the meanings of their lexical constituents are context invariant. *Qua* Wittgensteinian, I propose not to grant this assumption.’

Consider (3), an example that Travis has recently explored (Travis 2001). Suppose, for *reductio*, that (3) is compositional. (We ignore the demonstrative.) Presumably that would

3. This ink is blue

require the meaning of ‘blue’ to be context independent. A natural Platonist story would be that (barring idioms like ‘feeling blue’), 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 *the property of 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 doing *that* [sic], one might say any of many things” (Travis, 2001, p.197). Sometimes it’s how the ink looks in the bottle that decides whether (3) it’s true that 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 neither explication nor justification is provided. So semantics and metaphysics take in one another’s wash; the myth of the 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 content (Travis, 2001, pp.197ff).

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. For example: it appears there are aspects of language (/thought) that compositionality is essential to explaining. Systematicity and productivity are the classical candidates (for lots of discussion, see Fodor and Lepore, 2002). 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 given a context is the very burden of his plait. 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; deixis 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 appraised 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 this '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. 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' can determine what an utterance means. And so indeed it does, according to us. But we'll argue, on one hand, that this can't be metaphysical determination; and, on the other, that what we're about to say about the resolution of ambiguity holds, *mutatis mutandis*, for 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. Or rather: We think that they don't. 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 about what Groucho had in mind; that's the joke, don't you see. (There is *no* joke that philosophical analysis can't spoil.))

There are two, 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, as a 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 the contribution of the background is merely evidential. There could, no doubt, 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 another 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 iff (he in his pajamas) (shot an elephant) or whether it's true iff (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*.

There is, in short, 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*. Or, to put it the other way around, 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 be so if disambiguation supervenes on the speaker's intentions; for better or worse, one's intentions are just the sort of things to which one's access is typically privileged. 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?

If, in short, disambiguations supervene on facts about the background, then it would seem that one's disambiguating interpretations always might be inferences from one's knowledge of such facts. But Groucho's access to the truth conditions of his utterance *isn't* inferred from what he knows about the background; it *isn't* inferred at all. That's all just 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. Unless disambiguation is atypical of interpretation at large, the morals would seem to be:

- 3 The speaker's access to the interpretation of his utterance is epistemically privileged.
- 4 Nothing about the background of an utterance is metaphysically constitutive of its interpretation.
- 5 The function of background knowledge in interpretation is (only) to provide premises for the hearer's inferences about the speaker's intentions.

If, in short, the resolution of ambiguity is typical of interpretation at large, 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, Gareth 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* (1982, p.68).

One might indeed be lead to think so. 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, actually, 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*. 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 forwards*. So, then, how you should interpret an utterance of 'port' would seem to depend on what you assume he knows about (how much you assume he knows about) this piece of nautical jargon. The speaker may know this or he may not. 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" (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 forwards' rule; indeed, it's true even if the helmsman knows that the skipper doesn't know about the 'facing forwards' rule; presumably that's because the helmsman and the skipper must both defer to expert users of the term.

Well, maybe so; our intuitions 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. According to the Cartesian (against whom the question must not, after all, be begged) 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. The speaker is privileged in respect of what he means, but not in respect of how to say what he means.

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. English is, as we're forever being reminded, 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. 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 what you intended to say, is always (as philosophers put it) 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 his freedom to *think about* his left ear in any context that he chooses to do so; for, according to Cartesians, it's what a speaker is thinking of when he demonstrates that determines the object that is demonstrated. It is of course true notwithstanding, 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 chooses, utter 'That ear is on fire', thereby demonstrating my *left* ear (and thereby saying something patently untrue). What can't reasonably do is complain about my demonstration having been misconstrued.)

To be sure, all this turns on mere appeals to intuition. It's unclear to us how we could proceed otherwise, since the crucial test of a semantic theory is (we suppose) that it assigns semantic interpretations to the expressions of L in a way that speaker/hearers of L find generally intuitive. If a putative semantics for English says that the interpretation of 'Nixon is dead' is that Fidel lives, there is *prima facie*, something seriously wrong with it. We do admit, however, that the intuitions we've relied on may be deemed idiosyncratic by other English speakers; even, perhaps, those in a state of reflective equilibrium. 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 (1978, p.355).

We aren't deeply moved by this. It seems Kaplan is offering the point about what interpretation is natural for the linguistically competent observer as an argument that Groucho has (inadvertently) demonstrated the Agnew picture. But why prefer that 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 'He's the best philosopher,' couldn't 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?

The moral, in our view, is that it's a mistake to read such cases as showing that, metaphysically speaking, the bearer of content is the utterance-in-context. Rather, according to us, they illustrate a quite general perplexity that one feels when faced with what may be solecism; i.e., when a speaker's verbal gesture is inappropriate to what he pretty clearly intended to communicate: Is the interpreter to go by what the speaker actually uttered or by what, in the circumstances, he must have meant to say by uttering it? Compare: 'He said 'infer' but he must have meant 'imply';' here too there are conflicting intuitions about which interpretation of the utterance is right. And, since there's no principle that applies in the general case, it's unsurprising that such intuitions are often labile and very sensitive to the particulars of the example. Still, according to us, even though the context is irresistibly misleading as to the speaker's intentions, if the speaker's intention is given, so too is the interpretation of his utterance. Epistemology is one thing, metaphysics is another. 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 first point is implicit in what preceded. The usual way of thinking about 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, this 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 interpretation_N, which is something that symbols *have* (indeed, something that they have essentially). Prevailing philosophical opinion is that what *gets interpreted*, and what *has an interpretation* are the very same things: expressions in a natural (*a fortiori*, public) language. Indeed, interpretation_N arises from the process of interpretation_V. That being so, the metaphysics of content and the epistemology of 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 story 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. What *get* interpretation_V are utterances; what *have* interpretation_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, and then we're through. Most of the philosophical discussion of the mind in the last several decades has, and quite self-consciously, viewed the issues from the 'second person' perspective. This is perfectly reasonable if what you're primarily worried about is the epistemology of mental ascription. And it's perfectly reasonable to be primarily worried about the epistemology of mental ascription if the refutation of skepticism is among your long-term goals. *Prima facie* anyhow, the epistemological situation of the ascriber of a mental state is typically quite strikingly different from that of the ascribee; and it's generally supposed that it's only for the former that skeptical issues arise. That's because second person mental ascriptions are plausibly *inferential* in the normal case. (Even logical behaviorists think this; what makes their view special is just that they hold that the inferences in question are more or less demonstrative.) Whereas, in the case of first-person ascriptions of (conscious) mental states, inferential models are wildly implausible on the face of it. On these sorts of assumptions, there really is an asymmetry built into the epistemology of the mental, and the view from the first

person's perspective really is quite different from the second person's. We have no quarrel with this understanding of the epistemological landscape. Indeed, we think something of this sort has to be true if we're to suppose that mental state ascriptions have truth values at all.

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 surprised. But what's out of the question is that your mental states supervene on your 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; metaphysical states never are. That's why you don't have a chance of deriving your metaphysics from your epistemology, however hard you try.

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, pp. 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. "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, pp. 221-253
- _____, 'Afterthoughts,' in *Themes from Kaplan*. Oxford: Oxford University Press, J. Almog, J. Perry, and H. Wettstein 1989, pp. 567-614.

Loar, Brian, *Mind and Meaning*, Cambridge University Press. 1981.
Searle, John, 'The Background of Meaning', in *Speech Act Theory and Pragmatics*, eds.,
Searle, J., F. Kiefer and M. Bierwisch (Reidel: Dordrecht), 1980, pp. 221-232.
Travis, Charles, *Unshadowed Thought* (Cambridge MA: Harvard University Press,
2001).