Comments on Lasersohn

John MacFarlane

September 29, 2006

I'll begin by saying a bit about Lasersohn's framework for relativist semantics and how it compares to the one I've been recommending. I'll focus in particular on what is required semantically in order to secure genuine disagreement among the participants in a conversation.

I'll then say a bit about Lasersohn's notion of a "stance" and the complication he thinks it requires in the semantics of propositional attitude verbs.

Finally, I'll turn to the neat new arguments he has given against indexicalism and for relativist semantics.

Indices, Contexts, and Disagreement

Lasersohn and I agree on adding an index: for him it is a judge; I prefer something more abstract (a standard of taste). But I won't argue that here, and I'll adopt Lasersohn's talk of a "judge" standard for purposes of this talk.

We agree in thinking of the semantic contents expressed by sentences in context as having truth values relative not just to a world (and maybe time), but to this extra "judge" index.

Lasersohn says that two contents are "contradictory" just in case there is no world and judge relative to which both are true. In this sense, the contents expressed by "roller coasters are fun" and "roller coasters are not fun" (on his theory) are contradictory.

He then says: "This, in a nutshell, is the explanation for 'faultless disagreement.' In a system where context plays a role not just in assigning contents to sentences, but also in assigning truth values to contents, and where contexts are assumed to provide an individual on whose tastes the truth values of sentences may depend, it is possible for two sentences to express contradictory contents, yet each be true in its context."

But I think this can't suffice as an account of "faultless disagreement," because the fact that two speakers assert contradictory contents does not imply that they are genuinely disagreeing.

The point should be familiar to those who saw my talk here last year ("Nonindexical Contextualism") or my talk at the Pacific APA ("Relativism and Disagreement"). To see the point, just consider Kaplanian contents, which have truth values relative to world-time pairs.

- At 8 am, Abe says "The sun is up," thereby asserting a "tensed" content p that is true at $<@, 8am>$ and false at $<@, 11pm>$.

- At 11 pm, Ben says "The sun is not up," thereby asserting a "tensed" content that is contradictory to p. (That is, there is no time t and world w such that both these contents are true at $<t, w>$.)

- Adopting some useful terminology from Perry, we might say that although neither's assertion is *about* any particular time (the contents are both "time-neutral"), Abe's *concerns* 8 am, while Ben's *concerns* 11 pm.

- While Abe and Ben have asserted contradictory contents, they have not in any real sense disagreed with each other. Indeed, they can both agree that the other's assertion is "accurate" or "true given its context of use."

- No one should have any temptation to class this as a case of "faultless disagreement," because there is not any disagreement.

The point generalizes to all kinds of contents, even standard possible-worlds propositions. Such propositions are not *about* any particular possible world, though an assertion made using one *concerns* the world of utterance, and not other worlds. Assertions of contradictory propositions at different possible worlds do not "disagree," since they concern different worlds.

Thus, the fact that on Lasersohn's theory, an assertion of "roller coasters are fun" and an assertion of "roller coasters are not fun" will have contradictory contents, does not (by itself) imply that the two assertions disagree in any real sense. To ensure that, we need to move beyond discussion of contents and indices and talk about how the indices are "initialized" by features of context.

To make this clear, let's consider a variant of Lasersohn's theory that would agree with it in what it says about contents, but not imply any real disagreement between speakers who assert "roller coasters are fun" and "roller coasters are not fun," respectively.

This variant theory would treat the judge index exactly the way the world and time index are usually treated. We would say that an assertion should be regarded as objectively correct (accurate) just in case its content is true relative to the world in which the assertion is made, the time at which it is made, and the speaker as judge. On this theory, there would be an *assessment-independent* standard for evaluating assertions. Speakers who asserted contradictory contents could regard each others' assertions as correct or accurate. They would no more take themselves to be disagreeing than Abe and Ben would, in the example above.

This theory would be "relativist" at the level of contents. But, because it doesn't vindicate intuitions of "faultless disagreement," and because structurally it is no different from (say) Kaplan's theory in demonstratives, it seems better to think of it as a form of contextualism—what I last year called "nonindexical contextualism."

So here's a theory that does not differ from Lasersohn's in what it says about contents, yet does not secure genuine disagreement about matters of taste. How does Lasersohn's theory differ from it?

It differs by allowing that, in assessing others' assertions for objective correctness or "accuracy," we can consider the content asserted relative to a judge that is *not* the speaker. On Lasersohn's view, it is legitimate to adopt an "autocentric" stance in assessing an assertion (taking the judge to be oneself), or to adopt an "exocentric" stance (taking the judge to be some third party, usually different from both the speaker and oneself). This theory does vindicate intuitions of disagreement, because from an autocentric point of view it won't be the case that both assertions are objectively correct or "accurate."

It seems to me that these are two different *semantic* theories, in the sense of being two different theories of the meanings of predicates of personal taste. In the framework I've been developing, the difference between them can be seen easily from their different definitions of "truth at a context of use and context of assessment":

- *Variant Theory:* S is true as used at C_U and assessed from C_A iff the content of S at C_U is true at $< w, t, j >$, where $w =$ the world of C_U, $t =$ the time of C_U, and $j =$ the speaker of C_U.

- *Properly Relativist Theory:* S is true as used at C_U and assessed from C_A iff the content of S at C_U is true at $< w, t, j >$, where $w =$ the world of C_U, $t =$ the time of C_U, and $j =$ the assesser at C_A.

Lasersohn's approach is different. Instead of distinguishing two different "concrete" contexts, the concrete situation in which the sentence is used and that in which a use of it is being assessed, he sticks with Kaplan's formalism with one context, but loosens the connection between this context and the concrete situations of utterer and assesser. Thus, on Lasersohn's view, some features of the context will be taken from the utterer's situation (world), others from the assesser's (judge).

In some sense, this is just another way of representing the same theory. But I want to suggest that my approach, which goes for two "concrete" contexts instead of one abstract one, makes it easier to *see* how theories of the "nonindexical contextualist" sort differ from genuinely relativist ones. Briefly: a semantic theory is relativist if and only if it allows that truth (of sentences or propositions) can vary with the context of assessment.

Stances

A bit, briefly, about Lasersohn's distinction between "autocentric" and "exocentric" stances. Although my framework is flexible enough to cover both (to say that the judge is determined by the context of assessment is not to say that it must be the assessor), I don't see a compelling reason to allow the judge to differ from the assessor. It seems to me that the data suggestive of exocentric stances can be handled better in other ways.

Since time is short, let me focus on a paradigm example motivating talk of an exocentric stance.

> Mary: How did Bill like the rides?
> John: Well, the merry-go-round was fun, but the water slide was a little too scary.

Lasersohn says that John is here asserting the content *the merry-go-round was fun*, but taking up an exocentric stance, and that assessors will have a "tendency" to take up the same stance in evaluating it (taking the judge to be Bill).

Lasersohn also seems to assume that John's assertion expresses a belief with the same content, so that we can say "John believes that the merry-go-round was fun." This forces him to say that belief attributions are really three-place: they relate a believer, a content, and a judge (which may vary depending on the believer's "stance"). He draws an interesting distinction here with "consider," which he thinks is two-place, implying the autocentric stance.

I find the idea that "believe" is covertly three-place to be pretty implausible:

1. It seems odd to me to describe the case above by saying, "John believes that the merry-go-round was fun." Much more natural: "John believes that the merry-go-round was fun for Bill."

2. If "believe" had a third argument-place for a judge, shouldn't we expect to find it sometimes filled explicitly, or bound by a quantifier? Yet we don't. If Lasersohn is right, then there should be a content expressed by $\exists x(believes(John,\ \hat{}\ fun(the\text{-}merry\text{-}go\text{-}round,\ x)$. Indeed, this content would be just the thing to use to attribute belief when we know what proposition John endorsed but don't know what stance he was taking towards it. But how would you say this in English?

Moreover, if John has asserted the content $\hat{}\ fun(the\text{-}merry\text{-}go\text{-}round)$, then on Lasersohn's view it should be legitimate for others to take up an autocentric stance towards it and say, "John is wrong. Merry-go-rounds are never fun." But this just looks like a misunderstanding: John could justly say, "What I meant was that the merry-go-round was fun for Bill. I found it quite dull myself."

It seems to me that the data are best accounted for by saying that John is using the sentence "the merry-go-round was fun" to assert the content *^[fun for Bill](the-merry-go-round)*. He's relying on his listeners to see that he doesn't mean to assert the official semantic content of the sentence.

Lasersohn can't really object that there's no strictly compositional route from the sentence to the assertoric content, because he has to accept some degree of "free enrichment" of assertoric content himself. "Everyone did something fun" would naturally be used to assert that everyone x did something that was fun for x. But this reading won't be generated compositionally on Lasersohn's theory.

"Considers"

Here's how I understand Lasersohn's first argument against indexicalism:

1. "John considers the licorice to be tasty" has (only) a *de se* reading with respect to the judge. (It cannot report John's *exocentric* taking-true of the proposition that licorice is tasty.)

2. This reading is not captured by the indexicalist analysis: *considers(John, ^tasty(the-licorice, John))*. (And it would be implausible to attribute the *de se* reading to implicature, given that one can't even get the non-*de se* reading.)

3. Lasersohn's account, by contrast, does capture the *de se* nature of this claim. To stand in the "considers" relation to a content, Lasersohn claims, is to assess that content as true from an autocentric stance—that is, to take it to be true relative to oneself as judge.

4. The most obvious moves the indexicalist could make to solve this problem involve "significant concessions to the relativist analysis," e.g. the use of centered-worlds propositions as contents.

5. If the indexicalist says that the hidden experiencer argument of "tasty" "must be interpreted like a *de se* pronoun dependent on the matrix subject," this would be a case of control of a non-subject argument, which is not expected on current syntactic theories.

I think this is a neat argument against indexicalism. But it should be clear from what I've said already that it doesn't get you all the way to relativism. It motivates countenancing belief contents whose truth is relative to a judge, but (as discussed earlier) this is compatible not just with a real relativist position but with a kind of "nonindexical contextualism" that Lasersohn rejects.

This should be clear from Lasersohn's observation that Lewisian centered-worlds contents are formally just like his relativist contents. For Lewis was clearly not

defending any kind of view that would vindicate inuitions of "faultless disagreement." On the Lewisian view, as I understand it, a belief with a centered-worlds content is objectively correct, or accurate, just in case that content is true relative to the believer. Autocentric assessments of others' de se beliefs would, I suspect, have struck Lewis as just wrongheaded.

Thus I believe it is misleading to characterize the "two most central features of the relativist analysis" as constancy of content and relativity of truth of contents to a judge index, as Lasersohn does. These features can be shared by nonindexical forms of contextualism. The central feature of a relativist semantics, I would say, is variation of truth with the context of assessment.

"Recognizes"

Lasersohn's second and third arguments concern factive contexts. He divides them up under the heads of "factive predicates" and "truth-evaluative adverbs," but the arguments are more or less the same in each case. I want to focus on the constructions he uses as examples, "recognizes" and "correctly believes," and highlight a couple of aspects of these constructions that might merit further consideration. First, the argument with "recognizes."

Someone asserting

(L) John recognizes that licorice is tasty.

would be naturally understood to have committed herself to both of the following:

(a) Licorice is tasty to John.
(b) Licorice is tasty to the speaker.

This prediction falls right out of the relativist analysis. But it is at least very difficult to get it on an indexicalist account, which forces the judge into the *content* of the attitude attribution. If the thing John is being said to recognize is *that licorice is tasty [to John]*, then the factiveness of "recognize" commits the speaker only to licorice's being tasty to John, not to its being tasty to her.

A more sophisticated indexicalist might say that the thing John is being said to recognize is that licorice is tasty to a *group* of people that includes both John and the speaker. This would explain why the speaker too is presumed to like the taste of licorice.

Against this, Lasersohn argues that John might think he's the only person who believes licorice to be tasty, and the speaker might know this. In this case, he

says, it would still be felicitous and true for a speaker to say that John recognizes that licorice is tasty. And the indexicalist can't explain this.

But is it really so clear in this case that the speaker can felicitously and truly say that John recognizes that licorice is tasty? "Recognizes" is not just a factive verb; it is an epistemic verb. To say that John recognizes something is not just to say that he correctly believes it; plausibly, he must *know* it. Indeed, Timothy Williamson has argued that for all factive stative verbs V, if S Vs that p, then S knows that p; if this is right, then many of the factive constructions Lasersohn is considering here are *also* epistemic constructions. I want to suggest that the distinction between merely factive and epistemic contexts might be important here.

With this in mind, let's reconsider Lasersohn's claim. If John mistakenly thinks he's the only person who likes the taste of licorice, and if he forms his belief that licorice is tasty simply on the basis of his liking licorice, can the speaker truly say that John *knows* that licorice is tasty?

Here's an argument that she can't. In order to count John's belief that p as knowledge, the speaker must think not only that p is true, but that the fact that John has a true belief in this case is not just an accident. This might be cashed out in various ways, but no matter how it is cashed out, it is going to be problematic for a relativist:

- An old-fashioned way of cashing out the non-accidentality requirement is to require that John be *justified* in believing what he does. Well, in this case John has what he considers a justification: he likes licorice, therefore, licorice is tasty. But can *the speaker* regard this as a justification? After all, from her point of view, the truth of "John likes licorice" is very weak evidential support for the truth of "licorice is tasty."

- An externalist might talk instead about the *reliability* of the connection between John's believing that licorice is tasty and the truth of this belief. But certainly there's not going to be any reliable connection between John's believing that licorice is tasty and the truth of this belief (as the speaker assesses it autocentrically), unless there's a strong correlation between what John likes and what the speaker likes.

- One nice way of stating the nonaccidentality requirement, due to Nozick and recently revived by my colleague Sherri Roush, is that the subject's belief should "track" the truth: If it had been false that licorice were tasty, John would not have believed this; and if things had been different in various ways but licorice had still been tasty, John would still have believed this. Certainly the speaker has no basis for accepting these counterfactuals—unless, again, there is some reason to think that John's tastes are strongly correlated with hers.

If these considerations are correct, then on the relativist account, it's only going to be felicitous to assert that John recognizes that licorice is tasty when it is understood that John's tastes and the speaker's are the same.

But if we limit ourselves to these cases, the indexicalist will have a much easier time! For in contexts where it's presupposed that the speaker's tastes and John's are likely to be the same, an assertion with the content *licorice tastes good to John* will tend to convey that licorice tastes good to the speaker, too.

There is much more to be said here, but this should suffice to make the point that the distinction between epistemic and non-epistemic factive contexts may be highly relevant to the kind of argument Lasersohn is making.

"Correctly believes"

Lasersohn's "truth-evaluative adverbs"—"correctly", "accurately", and their negative forms—are not epistemic. So there is no problem saying that John correctly believes that licorice is tasty, even when there is no reliable link between his believing this and its truth. In these cases, however, the indexicalist has options that Lasersohn does not consider, and that are ruled out in the epistemic cases. For example, it may be that the speaker is attributing a belief with the content ˆ*tastes-good(licorice, those who are qualified to judge)*—specifying the relevant group through a property rather than extensionally. (It's less important what this property is, exactly—it could even be "those who have tasted both licorice and wormwood," or "those who have tasted the same things I have.")

This is something John could perfectly well believe even if he thought he was the only one who found licorice tasty. For he might believe that he is the only person qualified to judge. In this case, it would be perfectly fine for the speaker to say that John correctly believes that licorice is tasty (construing the content as just described). If, in the context, it is presupposed that the speaker is qualified to judge, the speaker will be understood as thereby committed to the claim that licorice tastes good to her.

Note that this response is not available in the case of "recognizes," because the speaker should not say that John's belief counts as knowledge in this case, given that John's belief is based on other false beliefs. (John just got lucky: even if the others qualified to judge had *not* liked the taste of licorice, he still would have believed that licorice was tasty.)

One final complication. It is not at all obvious to me that "correctly" and "accurately" *are* truth-evaluative adverbs, in Lasersohn's sense. For assessments of the correctness or accuracy of an assertion can come apart from autocentric assessments of the truth of its content. For example, to return to our temporalist example from the beginning, Ben will take Abe's assertion to be "correct" and "accurate" even though he takes its *content*—that the sun is up—to be (now)

false. Similarly, we will assess a counterfactual utterance of "Dodos are alive in 2006" as "accurate" if made in a world where Dodos aren't extinct, even though we take the content of the assertion—the proposition that Dodos are alive in 2006—to be false.

Since Lasersohn's argument depends on the assumption that "correctly believes" and "accurately believes" are factive, and this assumption is only true "up to approximation" (that is, only if we ignore these "corner cases"), the argument needs some adjustment. One option would be to use the construction "truly believes", but in idiomatic English this is most naturally interpreted as "sincerely believes." So perhaps we should use

> John believes that licorice is tasty, and that is true.

Surely *this* commits the speaker both to licorice being tasty for John and for it being tasty for the speaker! But now we are getting very close to the way ordinary "faultless disagreement" arguments run. We are appealing to our audience's willingness to accept the inference:

> John believes that licorice is tasty.
> What John believes is true.
> So, licorice is tasty.

Or, in its more familiar contraposed form,

> John believes that licorice is tasty.
> Licorice isn't tasty.
> So, what John believes isn't true.

Of course, Lasersohn's intent was to reduce our dependence on this kind of argument. If these considerations are correct, it is not clear how much he has really done so.