

What Sort of Science is Semantics?

Jerry Fodor and Ernie Lepore

We aren't happy about semantics. In particular, we aren't happy about current accounts of the goals of theory construction in semantics, or about how, if at all, semantic theories are supposed to illuminate the workings of the mind. We know, more or less, what we think about these questions; but we find it increasingly difficult to understand what the semanticists themselves say. We suspect cross-purposes, so we want to put some of their views on the table, where they can be scrutinized and compared with our own.

Here's what we're going to do: We'll set out a rough account of what we think semantics is about and, in particular, of how it ought to relate to undertakings that are familiar from Cognitive Science. The upshot will be that semantics is, or ought to be, part of Cognitive Science. In particular, it's the part of cognitive science that specifies the properties of the mind and of the world in virtue of which language and thought are about the world. In the most familiar case, the semantic theory of a language L specifies the situations in the world that are the truth (/falsity) makers of L-expressions; it explains, for example, why the current location of a certain contextually determined feline can make true (/false) a token of the English expression 'The cat is on the mat'. (Since we assume that (declarative) sentences are used primarily to express thoughts, we might just as well have said that the goals of a semantics of English include explaining why the *thought* that the cat is on the mat is true in, and only in, circumstances where tokens of the English *sentence* 'The cat is on the mat' would be.)

We're quite aware, of course, that this view of semantics is not universally shared,¹ but we do think there is some consensus as to the following:

Semantics is (or ought to be) an empirical science (like botany, entomology, geology and so forth) rather than a formal science (like logic or mathematics). So, for example, a semantic theory of L ought to be 'naturalistic' in the same way that theories in biology, meteorology, entomology, botany, astronomy, geology, and neurology are generally thought to be; in particular, it ought to be responsive to contingent data concerning L-speakers' intuitions about the semantic properties of expressions of L.

What follows is mostly a discussion of the methodological implications of this assumption; We'll argue that they bear closely on the goals and practices of theory construction in semantics. That's because, so far at least, the only plausible account of how semantic theories could be naturalistic and empirical requires assuming that it is part of Cognitive Science. We want to make clear from the outset that we take this *not* to be a merely terminological claim about

1 For example, It was in the spirit of American Structural Linguistics (see, e.g., Harris 1951) to think of grammars as no more than perspicuous, (or elegant, (or just brief)) summaries of linguistic data. In philosophy, Scott Soames has robustly championed this sort of view for several decades (cf., e.g., his 1984). But, as we'll see, the question then arises what could justify holding that facts about informants' intuitions (or indeed, any other empirical data) are germane to the (dis)confirmation of linguistic theories. More on this presently.

what is, or should be, *called* ‘semantics’; or about exactly how, or exactly where, the borders between linguistics and psychology ought to be drawn. To the contrary, the present questions are highly substantive; if it’s right to think of semantics as a naturalistic, empirical science, then quite a lot of what linguists, psychologists and philosophers have said about the relation between thought and language can’t be true. So, at least, we’re about to try to convince you.

To begin with, it bears emphasis that if a theory, *any* theory, is to be responsive to a body of empirical data, it must provide – or at least provide for – an intelligible account of why those data are germane to the truth or falsity of that theory. If you wish to claim that coming down with a fever and little pink spots is reliably a symptom of measles, you are in debt for an explanation of why *those* symptoms indicate *that* disease. If you claim that observations of the precise time at which Mercury is eclipsed by the Sun are data relevant to the (dis)confirmation of theories about the geometry of space-time, you are to that extent, in debt for an explanation of why *are*. Likewise (to borrow an example from Fred Dretske (1988)) if you claim that the reading of your fuel gauge is relevant to estimates of how much fuel is left in the tank, you are in need of a story (what we’ll call a ‘bridge theory’) about *why* it’s relevant. (You can read all about that in such publications as *Fuel Gauges for Dummies* (available, no doubt, from Barnes and Noble or Amazon)). And, all else equal, the more articulate and plausible your bridge theory is, the more you are justified in claiming that your data really do (dis)confirm the theory that you say that they do. As Dretske observes, what is common to all such examples is that states of affairs that are observed ‘carry information’ about states of affairs that aren’t. In the fuel gauge case, we have a pretty good bridge theory explaining why this is so; which is just to say that we know how fuel gauges work. As Timothy Williamson remarks, “we need a theory about the relation between the value of the quantity [that we wish to measure] and the representations of it we record when we use our instruments... The scientific investigation [of the quantity] widens to include [the theory of the instruments of observation]” (2004). Meehl and Chronbach got this sort of point exactly right in their classic 1955 paper on the role that considerations of the ‘construct validity’ of test procedures play in the evaluation of empirical theories. “Suppose we fail to find a greater incidence of “homosexual signs” in the Rorschach records of paranoid patients. Which is more strongly disconfirmed – the Rorschach signs or the orthodox theory of paranoia? The negative finding shows the bridge between the two to be undependable, but this is all we can say. The bridge cannot be used unless one end is placed on solid ground. The investigator must decide which end it is best to relocate” (1955, p. 231).

All of this continues to apply when the ‘instrument of observation’ is an L-speaker, whose intuitions about the meaning, reference, propositional content, truth conditions, truth makers, and so forth of L-expressions are the empirical data constraining semantic theories of L. Roughly, (and allowing for differences in terminology), the consensus view is that the *semantic properties* of an L-expressions (i.e. the semantic properties that the intuitions of L-speakers are taken to reveal) are determined by the *semantic values* of L expressions, i.e. by the senses, propositions, functions, or whatever, that linguistic expressions denote or otherwise express. If anything like that is so, we need an (empirical, naturalistic) bridge theory that explains how the intuitions of L-speakers can carry information about which L-expressions have which semantic values. If it belongs to the semantics of English that where the cat is bears on the

truth/falsity of tokens of ‘The cat is on the mat’, and if the evidence that it does is that English speakers say that it does, we need a story that makes it plausible that such intuitions are reliable. That is precisely the sort of question that a bridge theory for semantics is required to answer. At the very least, if we don’t have such a theory now, we will have to have one by and by.

There’s nothing about this that is peculiar to semantics. What makes the L- speaker’s *soi-disant* ‘intuitions of well-formedness’ germane to (dis)confirming theories about the syntax of expressions in L? There is a pretty general consensus that English-speakers have the intuition that ‘The cat is on the mat’ is well-formed; and that the fact that they do is part of the evidence for a linguist’s claim that it is well-formed. It’s our impression that even linguists who are neutral (or skeptical) about the ‘psychological reality’ of syntax are generally party to this consensus. Accordingly, the questions we’re raising are: *Who says that the informant’s ‘well-formedness intuitions’ are, in fact, intuitions of well-formedness? And, even assuming that they are, who gets to say we are justified in taking the informant intuitions of well-formedness to be reliable?* It is, after all, a methodological truism that being good at doing X doesn’t, in and of itself, justify one’s claims about how one does X. The informant gets to say what his intuitions are; but not what they are intuitions of, or whether they are true. These are typical questions about the construct validity of intuitive data in the evaluation of grammars. If there were evidence that having good linguistic intuitions about L is a necessary condition for being a competent speaker of L, that might justify taking the intuitions of competent L-speakers to constrain linguistic theories of L. But, to our knowledge, there is no such evidence. To the contrary, there are data suggesting, for example, that children quite generally achieve considerable linguistic competence before they come to have linguistic intuitions, reliable or otherwise.

We take the issues about the construct validity of intuitive data to be crucial in meta-theoretic discussions of linguistics, semantics included. For example, it is perfectly possible that semanticists are *just wrong* when they say that modal intuitions bear on the (dis)confirmation of their theories. English speakers generally agree that you can’t have killed anyone who isn’t dead; and that they do is widely taken to be evidence that ‘kill’ means CAUSE TO DIE. Likewise, the modal intuition that there is no possible world in which H_2O is anything but water is thought to argue that it’s a *linguistic* truth that water is H_2O . But such claims want justifications. Maybe what explains the informant’s intuition that ‘if killed then dead’ is necessary is *not* the meaning of ‘kill’ but the metaphysics of killing. In that case, though it is *necessary* that *killed entails dead*. Nothing follows about whether it is analytic. Rather, what’s wrong with someone who doesn’t understand that ‘killed \rightarrow dead’ is necessary isn’t that he doesn’t know what ‘kill’ means; it’s that he doesn’t know *what killing is*. In fact, many thoroughly reputable philosophers have been skeptical about ‘truth in virtue of meaning alone’; that is, they have denied that modal intuitions are intuitions of semantic relations. Quine and Goodman are, surely, about as respectable as philosophers get; but both held that modal intuitions are neither semantic nor metaphysical but epistemological. Quine took them to be intuitions of theoretical ‘centrality’ (Quine 1951), and Goodman thought them to be intuitions of theoretical ‘entrenchment’ (Goodman 1979).

Our point in all this is *not* to revisit old issues about whether modality is a linguistic property or whether there is such a thing as analyticity. Our point is just that, if you claim that intuitions

of modality or of analyticity, are germane to the confirmation of semantic theories, you are in need of a bridge theory to explain why it's right to do so; it can't be 'true by stipulation', or 'true by definition', that a datum counts for (or against) a theory; not, anyhow, if the theory purports to be *empirical*.

We're so sure that getting straight about that is central to understanding how the (dis)confirmation of empirical theories works, that we propose to squander a few more paragraphs rubbing it in; then we'll consider how it applies to semantic theories in particular.

In the bad old days of Operationalism, it was widely held that empirical theories should *entail* their data. That seemed plausible because it was likewise widely held that empirical theories must include 'coordinating definitions' which connect the 'theoretical vocabulary' to the vocabulary in which the data are. For example, the geologist's claim that diamond is harder than chalk is confirmed by his observation that diamonds scratch chalk, and not vice versa. That it follows from the 'operational definition' of 'x is harder than y', which invokes the relation 'x scratches y' (it's something like: 'x is harder than y iff, all else equal, x scratches y and not vice versa.'). A more sophisticated Operationalism might hold that the scratch test is (not an operational definition of hardness but) a 'criterion' for hardness attribution. (This sort of difference between kinds of operationalist once seemed very important; but it doesn't any more.)

Operationalism was, of course, useful to epistemologists, for whom the refutation of skepticism was what finding the Grail was to King Arthur. The skeptical worry is: because of observational error and the like, *even those theories that best conform to the data may not be true*. If, however, the connection between the theory and the data is somehow definitional, this worry is incoherent: If 'is harder than' *means* SCRATCHES, the geologist's inference from the latter to the former is ipso facto philosophically licensed. (Not entirely surprisingly, the view that geologists require work-permits from philosophers is more pervasive among philosophers than among geologists.) Likewise in the case of semantics: if a grammar is *by definition* a systemization of the L-speaker's intuitions, then linguists don't, after all, require a bridge theory to justify appealing to such intuitions as evidence for (or against) the empirical adequacy of grammars. All they need appeal to is the definition.

But it turns out (see the previous discussion) that the connection between an empirical theory and its confirmation base are generally *not* sustained by definitions, operational or otherwise. Rather, (as Duhem famously emphasized) the confirmation of an empirical theory typically recruits a complex of highly substantive assumptions, including highly substantive assumptions about how the instruments of observation work. These may be drawn from practically any source that seems reliable, including: theories in other sciences, results in engineering, and even thought experiments (to say nothing of brute common sense). It was a matter of dispute among early astronomers what, exactly, telescopic evidence is evidence of canals on Mars? Scratches on the lens?). This dispute wasn't settled by consulting the definition of 'telescope'; nor could it have been. Meehl and Chronbach were absolutely right.

So Operationalism died. So now what? In particular, now what about the relations between the intuitions of L-speakers and the semantics of L? We've remarked that many linguists hold that the goal of semantics is to capture the 'modal intuitions' that informants have about ex-

pressions in their language; in particular, intuitions about which expressions are entailed by which others. Many philosophers say much the same about ‘conceptual analysis’. Thus Frank Jackson (1998) writes: “Our subject [in philosophical analysis] is really the elucidation of the possible situations covered by the *words* [sic] we use to ask our questions – concerning free action, knowledge, and the relation between the physical and the psychological and whatever...” (p. 33). If that’s true, not much can turn on the difference between what linguists do when they do semantics and what philosophers do when they analyze concepts.

But though they frequently appeal to the semantic values of L-expressions (e.g. to the satisfaction conditions of L-sentences, or the meanings or the ‘uses’ of L-words) to explain the modal intuitions of L-speakers, philosophers and linguists have both been remarkably cavalier about why it’s alright for them to do so. After all, semantic theories are supposed to be naturalistic explanations of semantic intuitions; and the usual case is that empirical theories explain their data by postulating causal interactions between things in the world and instruments of observation. But the semantic values that semantic theories typically postulate (*propositions, meanings, functions, sets of worlds* and the like) *don’t have causal powers; they don’t make things happen*. A fortiori, they don’t cause speakers to have grammatical intuitions. Why, then, is it all right to assume that the grammatical intuitions of speakers bear on the evaluation of grammatical theories; and why don’t philosophers and linguists worry about that?

To put the point in a nutshell, lacking a credible bridge theory, the relation between semantics and speaker-intuitions does *not* seem to conform to familiar patterns of empirical explanation. That’s because, on the one hand, the sorts of semantic values that semantic theories take L-expressions to have, aren’t *causal* properties; a fortiori, they can’t figure in causal explanations of linguistic intuitions. So, if the grammatical intuitions of competent speakers are to constitute data for semantic theories to explain, we need some sort of account of how the grammatical intuitions arise from linguistic competence (or, conceivably, vice versa); and, unless grammars are psychologically real, *we don’t have one*. Why don’t linguists and philosophers worry about this? Why don’t they spend long days and nights trying to construct a bridge theory about how it could be that semantic facts explain linguistic intuitions?

Perhaps, it’s because they think they already know how. True enough, propositions, meanings and the like don’t make things happen, but *propositional attitudes* do. The proposition that the cat is on the mat is causally inert, but John’s believing that the cat is on the mat can contribute causally to determining what else he believes and what he is disposed to do. That being so, maybe we don’t, after all, need a bridge theory for semantics in order to make it bona fide empirical. All we need assume is that L-speakers believe that L-expressions have the semantic values that a correct semantic theory for English says that they do (a kind of belief that L-speakers might be supposed to pick up when they learn to speak L). If a correct semantics for English says that the proposition that the cat is on the mat is the semantic value of the expression, ‘The cat is on the mat’, *it just follows* that that English speakers generally believe that the semantic value of ‘The cat is on the mat’ is the proposition that the cat is on the mat. And, unlike propositions, facts about what people believe *do* make things happen. For example it’s perfectly ok for John’s believing that the cat is on the mat to contribute causally to John’s saying that that’s where the cat is, should the question arise.

But, on second thought, that won't do. If it is a puzzle how propositions could cause *intuitions*, it is likewise a puzzle how propositions could cause *beliefs*. After all intuitions just *are* beliefs; having the intuition that 'The cat is on the mat' means that the cat is on the mat just is having the belief that 'The cat is on the mat' means that the cat is on the mat. If there's a puzzle about the second, there must be the same puzzle about the first,

Bear in mind that we're looking for a defense of the claim that speaker-intuitions are germane to the (dis)confirmation of semantic theories. Well, for those purposes, to take it for granted that the intuitions that speaker's pick up when they learn *L* *are* intuitions of meaning would be to argue in a circle. If we just stipulate that semantic theories are about the meanings of expressions, we can't also just stipulate that speaker intuitions bear on the confirmation of semantic theories. It's no use explaining the informant's *intuitions* about modality by advertising to his *beliefs* about modality since, if a bridge theory is required to explain either, then it is likewise required to explain the other. So we're back where we started; all this dog catches is its own tail.

That's one reason why mentalistic bridge theories so often strike one as question-begging (much as Quine warned that they would). But we want to emphasize that what's worrying us is *not* what worried Quine. As far as we can tell, Quine thought that explanations involving attributions of propositional attitudes are suspect in and of themselves. Propositional attitudes are creatures of darkness; you won't find them in first-class conceptual systems. By contrast, we have no quarrel with propositional attitudes *per se*; some of our best friends have them. What we're worried about is propositional attitudes *about things that lack causal powers*.

Why We Think Semantics had Better be a Branch of Cognitive Science

The proposal is simple and it's nothing new: To claim that semantics is a branch of cognitive science is to claim that the semantic properties of propositions of linguistic expressions are *mentally represented* by linguistically competent speakers, and that such representations are etiologically involved in the formation of a competent speakers linguistic intuitions. That 'cat' expresses the concept CAT (and hence, cats belong to the extension of that concept) is one of the things that English speakers know about English. Concepts are, of course, abstracta, so they have no causal powers. But mental representations of concepts are concrete, so they do. The tokening ('in your head' presumably) of a representation of a concept is like the tokening of an English sentence-type. Both such events have causes, and both make things happen. Our uttering 'Stop that!' may cause you to stop that; and, if it does, part of the story about why it does is that, qua English speakers, we are disposed to mentally represent tokens of 'Stop that!' as meaning STOP THAT.

If you want the sorts of things that semanticists talk about to be the sorts of things that do causal work, they had better not be abstract objects. Mere common sense takes us that far.

Cognitive science adds that, all else equal, our intending to utter a token of the ‘Stop that’-type involves my mentally representing the intended utterance *as* a token of the ‘Stop that’-type; and that, all else equal, your hearing an utterance of the ‘Stop that’-type disposes you to mentally represent the utterance as intended to be a token of the STOP THAT type; and your understanding why we said that we had that intention. Cognitive science adds the thesis that mental states and processes are causal chains of tokenings of propositional attitude types, and it adds that the links of such chains are tokens of mental representations.

Short form: Many linguists and philosophers think that the thesis that the semantic properties of linguistic forms are mentally represented by speakers is a sort of optional add-on in understanding semantic theories. They claim, in effect, that you can psychologize linguistics if you are so inclined, but that understanding how linguistic explanations work doesn’t require that you do. We think that’s just plain wrong. If linguistic data include speaker intuitions, a coherent bridge theory for linguistics requires that grammars be internally represented; otherwise it is simply a mystery *why* linguistic intuitions are data for semantic theories.

We’re aware, however, that we’ve been arguing quite a narrow case. For, suppose we’re right to claim that you can use intuitive data to (dis)confirm semantic theories if you assume that linguistic intuitions are effects that are caused by mental representations of grammars. That *wouldn’t* show that, even in principle, semantic theories are naturalistic or empirical; whether they are would then depend on whether theory of mental representation can themselves be naturalistic and empiricistic. Perhaps, it goes without saying that whether they can is not a question to which decisive answers are currently available. Representation, mental or otherwise, is itself a semantic notion; so no wonder that so many linguists and philosophers suspect explanations of semantic facts that depend on assumptions about mental representations are bound to end in circularity.

But, in fact, *grammars had better* be mentally represented whether or not mental representation can itself be naturalized. After all, nobody really knows what naturalization amounts to, or whether a respectable metaphysics demands it of respectable scientific theories. But, metaphysics to one side, it can’t be a *miracle* that speakers have intuitions about the grammaticality, modality, and so forth of expressions in their language; and it can’t be a miracle that, by and large, such intuitions are accurate. We don’t know how to explain such facts without assuming that informants mentally represent the grammar of their language, and that mental representations play a critical role in the etiology their intuitions. Nor, as far as we can tell, nor does anybody else. To repeat just once again, *questions about construct validity are not a priori, and they must not be begged.* That’s one of the reasons that science is so hard.

A Brief Appendix on What Chomsky Has Been Saying of Late²

We suppose it's clear that discussion in the body of this paper presupposes a Realist view of notions like 'mental content', 'intentional content', 'mental representation' and the like. After all, the main line of argument has been that you need such notions if you're to make sense of the traditional linguistic practice of taking informant intuitions as germane to the (dis)confirmation of grammars; whereas Chomsky's current view seems to be a sort of Operationalism about the cluster of notions that circle around. They come in as *façon de parler*. But they aren't to be taken with full scientific seriousness. In the long run, precision and disambiguation will require their replacement by a technical vocabulary that dispenses with notions like mental content and mental representations with one that is explicitly *not* intentional. "One may, if one likes, say that I-Language [roughly, whatever it is that the speaker knows about his first language in virtue of which he is able to speak it] ... but this [usage] merely points to the problems to be investigated and should not be misinterpreted as a substantive proposal". Now, we are elderly and feeble chaps, and shocks are bad for us; and we simply can't get it through our head that Chomsky has joined forces with the likes of Dennett, the Churchlands, Skinner and Quine at his most Positivist. In fact, irony of ironies, we had thought one of the things we had learned from Chomsky is that you just *can't* have a serious psychology (or linguistics) that isn't flagrantly committed to intentional notions, 'mental representation' being very much among them.

We're too old to go through that all again; but let us recall a famous point of Frege's: Ordinary, everyday, talk about the mind is intentional through and through in the (we should have thought) reasonably precise sense that substitution of co-referents fails to preserve truth in the scope of many of our ordinary, every day psychological terms: 'John believes that Cicero was fat' and 'John believes Tully was skinny' may both be true, even though Cicero and Tully were the same Roman. This does rather suggest that intentionality is rampant in our quotidian understanding of the mental. Moreover, there is a direct (if not mandatory) route from Intentional Realism to Realism about mental representation, the key idea of which is that, though Cicero is the very same person as Cicero, representing that person as Cicero is not the same as representing that person as Tully; so it is perfectly possible to do the first but not do the second. It bears remarking on that this pattern of explanation runs throughout the whole of the cognitive/perceptual psychology that we've actually got. 'Subjects see the thing as a duck' and 'Subjects see the thing as a rabbit' can both be true even if it's the very same thing both times. The standard view is, of course, that 'seeing ... as ...' is intentional.

It seems to us that faced with these sorts of considerations; Chomsky has two options, neither of which is digestible.

1. Rule out intentional psychological explanations by brute fiat and accept the consequent loss of explanatory power.
2. One of the things Chomsky has been saying of late is that there's no fundamental difference between what he says now and what he's always said. We think he's quite wrong about that; but the question is merely hermeneutic, and we don't propose to discuss it here.

That's the strategy that Skinner tried in regard to the mental; without, we think, much luck.

2. Try to reduce the intentional to the biological/neural (assuming that the biological neural vocabulary isn't itself intentional). Sometimes it seems to us that this is what Chomsky now has in mind.

So far, however, this strategy has met with no successes either. In fact, things seem to be going the other way. Neurological texts are up to the ears in functional/teleological explanations. And functional/teleological vocabulary is (as far as anyone knows) also up to its ears in intentionality. Neurological explanation rejoices in such explanations as 'because it functions to maintain metabolic stability' and 'because it sharpens the retinal image.' So far, there is no ground for supposing they are eliminable *salve* explanatory power. (Many people think that Darwin showed how to reduce teleology to 'selection for 'It turns out, however, that 'select for ...is intentional to. (See Fodor and Piatelli (2011).)

For better or worse, intentionality is the mark of the mental; there seems to be no way to avoid it; and what cannot be avoided, one just has to live with.

Acknowledgements

Professor Gerhard Preyer has been an inspiration (and, at least to Lepore, a great friend) for years now. His contributions to philosophy have been significant and far-reaching. It is an honor for us to be able to contribute to this volume in his well-warranted volume celebrating his career. The topic of our paper is apt given the important contributions Gerhard has made in in semantics and in cognitive science and theory of mind.

Bibliography

- Cronbach, Lee J. and Paul E. Meehl, "Construct validity in psychological tests", 1955, *Psychological Bulletin*, 52, 281–302.
- Dretske, Fred, *Explaining Behavior*, MIT Press, 1988.
- Fodor, J.A. and M. Piatelli, *What Darwin got wrong*, Picador, 2011.
- Goodman, N. *Fact, Fiction and Forecast*, Harvard University Press, 1979.
- Harris, Zellig *Methods in structural linguistics*. University of Chicago Press, 1951.
- Jackson, F., *From Metaphysics to Ethics: A Defence of Conceptual Analysis*, Oxford: Oxford University Press, 1998.
- Quine, W.V.O., "Two Dogmas of Empiricism," *The Philosophical Review*, 1951: 20–43.
- Soames, Scott. "Linguistics and Psychology." *Linguistics and Philosophy*, 7, 1984: 155–179.
- Williamson, T., 'Past the linguistic turn?', in B. Leiter, ed., *The Future for Philosophy*, Oxford: Oxford University Press, 2004: 106–128.

