

[forthcoming, *Nous*]

Can Empirical Theories of Semantic Competence Really Help Limn the Structure of Reality?

Steven Gross

University of Pennsylvania

I. Introduction

There is a long tradition of drawing metaphysical conclusions from investigations into language. This paper concerns one contemporary variation on this theme: the alleged ontological significance of cognitivist truth-theoretic accounts of semantic competence.

According to such accounts, human speakers' linguistic behavior is in part empirically explained by their cognizing a truth-theory. Such a theory consists of a finite number of axioms assigning semantic values to lexical items, a finite number of axioms assigning semantic values to complex expressions on the basis of their structure and the semantic values of their constituents, and a finite number of production schemata. The theory enables the derivation of truth-conditions for each sentence of the language: something of roughly the form 'S is true iff P'.¹ The claim that speakers stand in a cognitive relation to such theories is advanced, not as a conceptual analysis of semantic competence or understanding, but rather as an empirical hypothesis about human speakers in particular, one part of a broader empirical account of our linguistic competence and cognition generally. It therefore must mesh with the rest of our

theorizing in these areas and whatever relevant data from neighboring inquiries there may be. (For example, since it's hypothesized that 'S' in the schema above should be replaced by a certain sort of syntactic representation of the sentence, syntactic evidence can bear on the semantic theory and vice versa.) The precise nature of the cognitive relation a speaker is supposed to bear to a truth-theory is a matter of some dispute. I speak of "cognizing" (following Chomsky 1975, pp. 164-6) as a place-holder for whatever the posited relation or attitude might turn out to be, allowing that it may in some ways diverge at least from typical cases of such more familiar attitudes as belief and knowledge.²

Cognitivist truth-theoretic accounts of semantic competence are thus neo-Davidsonian: they follow Davidson in utilizing truth-theories; but whereas Davidson claims only that someone who knew a truth-theory for a language would understand it, these cognitivists maintain that a speaker's semantic competence in fact is explained at least in part by their bearing some cognitive relation to a truth-theory.³

¹ I bracket complications proposed to accommodate context-sensitivity.

² Some further terminological notes. (1) There are many projects that fall under the heading 'semantics.' I use the longer phrase 'an account of semantic competence' to ward off confusion of my subject with other semantic projects. This becomes particularly important below when I invoke Scott Soames' distinction between semantics and semantic competence. (2) To avoid prolixity, I often speak simply of accounts of semantic competence when it is clear I mean cognitivist truth-theoretic accounts of semantic competence. (3) I indicate the semantically relevant syntactic representation of a sentence by surrounding it with square brackets. (4) I reserve 'account' for the empirical explanation of the semantic aspect of linguistic behavior, which on the view at issue involves the ascription of a truth-theory to a speaker; I reserve 'theory' for the truth-theories themselves that such an account ascribes. (5) I shall sometimes refer to theorists who draw ontological conclusions from a cognitivist truth-theoretic account of semantic competence as ontologizing semanticists.

³ See Davidson 1967a. For a textbook presentation of cognitivist neo-Davidsonianism, see Larson and Segal 1995.

Theorists have drawn both positive and negative ontological conclusions from such accounts of semantic competence, conclusions concerning both what does and what doesn't exist. A well-known positive example stems from Davidson's proposed treatment of some adverbs as involving quantification over, and thus ontological commitment to, events: our competence with such sentences as 'Emma is walking slowly' on this view is explained in part by our cognizing a truth-theory that yields a theorem to the effect that the sentence is true iff there's an event that's done by Emma, is a walking, and is done slowly. Quantification over events is introduced to handle a variety of other constructions and phenomena as well.⁴ Similarly, depending on the details of the truth-theory ascribed, one might argue for the existence of possible worlds, properties, kinds, sets, etc.

Peter Ludlow's (1999) recent work on tense provides an example of drawing a negative ontological conclusion from a cognitivist account of semantic competence. Standard treatments posit quantification over past and future events and/or times. For example, according to such an account, a pluperfect sentence such as [Emma had walked] is true, relative to the time of utterance and a contextually given reference time, iff there's an event in which Emma walks that occurs prior to the contextually given reference time that itself occurs prior to the time of utterance. Ludlow's account avoids such

⁴ See Davidson 1967b and Larson and Segal 1995, chapter 12. For discussion of events generally in the semantics of English, albeit not in a cognitivist setting, see Parsons 1990. My rough characterization of the relevant t-theorem is closer to proposals utilizing "thematic-roles" than to the original Davidsonian proposal from which they stem. A less rough characterization would inter alia need to accommodate the attributive nature of 'slowly', perhaps by having the event be slow for a walking or for the contextually relevant kind of event it is. But since I focus on the philosophical implications of such accounts generally, I won't be fussy about linguistic details that would only prove distracting.

quantification by instead deploying a tensed meta-language.⁵ [Emma had walked], for example, is true, according to his proposal, iff the Interpreted Logical Form of [Emma is walking] was true before [some contextually given claim] was true.⁶ Ludlow discusses a wide variety of considerations in assessing these alternatives, including some suggestive developmental data. According to the standard account, though not according to Ludlow's alternative, competence with even simple tenses presupposes minimal grasp of temporal-ordering concepts such as EARLIER-THAN and LATER-THAN, or BEFORE and AFTER. But acquisition studies suggest that children exhibit competence with tense well before exhibiting competence with terms expressing temporal-ordering concepts.⁷ This illustrates, incidentally, how such theorizing, as part of an account of semantic *competence*, is intended to mesh with other empirical studies of cognition to yield testable predictions.

Now, these competing accounts of our semantic competence with tensed sentences differ greatly in their ontological upshot, according to Ludlow. If the standard account were correct, we would be committed to the existence of past and future events and/or times. His alternative account, on the other hand, supports a "presentist" conception of time, according to which, of times and events, only the present moment

⁵ The meta-language is thus context-sensitive. Cf. Ludlow 1999, pp. 63-4. For some discussion of context-sensitive truth-theories, see Gross forthcoming.

⁶ An Interpreted Logical Form (ILF) results from pairing the nodes of the semantically relevant syntactic representation with their semantic values as assigned by the truth-theory. The pluperfect example illustrates Ludlow's general strategy of handling complex tenses by relating the truth-conditions of simply-tensed sentences via temporal connectives. See Ludlow 1999, chapters 7 and 8.

⁷ Ludlow 1999, pp. 138-41. This evidence is not decisive: for instance, it might show only that temporal-ordering concepts are lexicalized only after competence with tense is secured. Ludlow might reply, however, that his suggestion that an individual's I-language is her Language of Thought precludes this possibility. See Ludlow 1999, pp. 22-6 and 165-9.

and events now occurring are real. Using McTaggart's (1908) labels, it is claimed that the standard account of tense supports the B-series conception of time, according to which all events and times—past, present, and future—exist and are equally real; whereas Ludlow's alternative supports the A-series conception, according to which there is a fundamental metaphysical difference between the present, on the one hand, and the past and the future, on the other, in that only the present—or, what exists in it—in fact is real, or exists.

Can a cognitivist truth-theoretic account of semantic competence really yield this sort of metaphysical result? One might counsel caution for a variety of reasons. Most dramatically, doubts concerning the very enterprises of semantics or of ontology—or of both—would temper any optimism that conclusions from the former might shed light on issues in the latter.⁸ I place such qualms to one side for the purposes of this paper. They are important and instructive, but there's also room for worry even if one maintains the legitimacy of ontological questions and the scientific bona fides of empirical theories of semantic competence. I put to one side as well challenges—including empirical challenges—to this particular approach to semantic competence: the cognitivist truth-theoretic approach could turn out to be on the wrong track, but I will assume enough explanatory work has been done to make it plausible, not just that there's a science to be had of semantic competence, but that it might take this general form.⁹ Of course,

⁸ Hilary Putnam, for instance, holds that ontological disputes owe their specious impression of intelligibility to the false presupposition that such expressions as 'exists' have only one sense or use. (For a discussion with references, see Gross 2002.) Quine, on the other hand, while countenancing ontology, famously questions the naturalistic acceptability of semantics.

⁹ I thus put to one side, for example, Chomsky's pessimism concerning "referential" semantics, and pragmatics-based challenges (e.g., Carston 2002) to such semantics playing a role in an account of linguistic competence. I put to one side as well challenges

objections to particular ontological conclusions might be raised by challenging the particular semantic treatments upon which they are based—without challenging either the enterprise of semantics as a whole or, more specifically, that an adequate account of semantic competence involves the attribution of a truth-theory. Thus Ludlow challenges arguments for the existence of past and future events grounded in the received treatment of tense; and thus one might challenge Ludlow’s presentism, if his alternative treatment is likewise open to criticism. My concern below, however, is the move itself from an account of semantic competence to ontology. Ludlow holds that accounting for competence with tense has metaphysical upshot no matter which treatment proves correct. It’s the basis for that claim—and parallel claims in other cases—that I wish to examine.

Prima facie it would seem possible in principle, according to the ontologizing semanticist, that there be two competing truth-theories such that the ascription of each equally well accommodates all relevant evidence except (say) the order of acquisition of some particular elements of the language. In this case, it would seem that adverting to a contingent feature of human cognitive development might settle some matter of metaphysical moment.¹⁰ But perhaps we should be wary of basing our ontology on evidence drawn from contingent features of human physiology, development, cognition,

(cf. Pietroski 2003) that allow that an account of semantic competence might utilize t-theories but question whether the semantic value represented by ‘t’ is truth.

¹⁰ For example, were there compelling responses to Ludlow’s objections to standard treatments of tense (or, as Ludlow 1999, p. 138, puts it, if “we decide to ignore” the other considerations he adduces), an ontologizing semanticist might allow her position on the reality of the past to hinge on the outcome of developmental research of the sort Ludlow finds suggestive. Cf. Ludlow 1999, p. 40, on the need for psychological evidence: “With adequate ingenuity, it may be possible to construct several possible semantic theories for a given natural language phenomenon (by hypothesis, theories with different metaphysical consequences).”

or evolution. Might it likewise be in principle possible that there be non-humans who speak a language broadly mutually intelligible with respect to ours, but whose pattern of acquisition differs in the relevant respect?¹¹ Would this commit them to a different ontology? Should it make us rethink ours? What metaphysical conclusions should an ontologizing semanticist draw who uncovers this difference in cognized truth-theories after decades of co-mingling among the species?¹² It's natural to think that, if the species could happily co-mingle, without any metaphysical differences coming to light through casual conversation, then there's something odd about attributing varying ontological commitments across the species and something even odder about drawing ontological conclusions from the inner workings of one language as opposed to the other. Such speculation might not bear much weight—perhaps no such case is in principle possible after all, at least nomologically.¹³ But it should at least give us pause, and encourage a closer examination of the arguments an ontologizing semanticist might deploy.

¹¹ I advert to non-human speakers because the ontologizing semanticist might draw her conclusions from phenomena found universally among human first languages—the need for an eventish truth-theory might be an example. The developmental data adverted to would then presumably also be species-wide. Note that phenomena found only within a particular human language (e.g., that it happens to treat the noun for a certain kind as mass while others treat it as count) can seem too “superficial” to be of much metaphysical moment. But perhaps even cross-linguistic phenomena should likewise strike us as “superficial” relative to the linguistic behavior and competence of possible non-humans. Empirical evidence might have overturned Quine's (1960, p. 8) topiary analogy with respect to humans, but perhaps something akin to it might hold with respect to all (nomologically) possible speakers.

¹² Suppose the two languages aren't broadly mutually intelligible, but we're capable of learning the alien tongue. Suppose further that the cognized truth-theory of our native tongue doesn't involve quantification over past and future events while that of the alien tongue does. Do our ontological commitments increase when we learn the alien tongue as a second language? Should this depend on the cognitive details of second language learning?

¹³ It would seem unrealistic to expect the species' cognitive differences to manifest themselves in just one way. But, for the case to be possible, what's needed is just that the manifestations not be such as to preclude mutual intelligibility broadly speaking.

In what follows, I distinguish two broad strategies for moving from attribution of a truth-theory to ontological conclusions and, within them, several sub-strategies. My discussion has three main goals. The first is to bring out, in distinguishing the variety of strategies available, that it's not as clear as it could be just how arguments from accounts of semantic competence to ontology might proceed.¹⁴ The second is to highlight, for each option mentioned, its particular commitments and correlative possible points of resistance. I will record some commitments towards which some resistance is familiar (such as that the relevant quantifiers be in fact ontologically committing), but also some commitments less frequently challenged in this setting—for example, that the ascribed truth-theory be believed, that it be true, and that it be revealing of logical form. It's because I focus on truth-theories as part of an account of semantic competence that the possibility of questioning these further commitments opens up. I suggest that it's unclear whether explaining the relevant aspects of linguistic behavior requires taking these commitments on—in part by appealing to the possibility that speakers are something like tacit, or subpersonal, fictionalists concerning the ascribed truth-theories. The third goal is to illustrate a more general point. The cognitivist emphasis and the concomitant integration with empirical linguistics and other areas of cognitive science can be characterized as a particular way of naturalizing this region of philosophy of language (cf. Ludlow 1997a), one that represents a significant shift from earlier work in truth-theoretic semantics. This shift arguably clarifies and helps legitimize the research program's aims and claims. More controversially, it might be held to broaden significantly the empirical base available to address the philosophical projects for which

¹⁴ Various people have said to me: “Obviously this sub-strategy is the one ontologizing semanticists intend.” But I've been told this about each of the sub-strategies.

reflection on language has sometimes been enlisted. It needs to be examined, however, whether, in becoming incorporated into empirical research projects, such theorizing retains (or achieves) the philosophical upshot—for example, for ontology—that it has in the past been claimed to have. As my tentative wording indicates, I certainly don't settle the question posed in my title, but only map out some possible areas of contention.

II. Quinean Ontological Commitment and Two Positive Strategies

It will prove useful for distinguishing strategies for moving from attribution of a truth-theory to ontological conclusions if I first rehearse a few elements of Quine's conception of ontological commitment. Quine's views, when not highlighted in the foreground, often lurk in the background of contemporary discussions of ontology. But it's not being suggested that ontologizing semanticists endorse—or ought to endorse—all aspects of Quine's approach, the details of which are not my concern here.¹⁵

We can sum up Quine's criterion of ontological commitment as follows:

a theory—or, set of beliefs—is committed to the existence of *k*'s iff it logically implies that there are *k*'s, and a person is committed to the existence of that to which her beliefs are committed.¹⁶

¹⁵ Certainly they don't share Quine's negative views on the scientific status of semantics. For Quine on ontological commitment, see, e.g., Quine 1948, 1960, pp. 238-4, and 1981. (Incidentally, the title of the present paper alludes to his (once well-known) remark about "limning the true and ultimate structure of reality." (Quine 1960, p. 221) To limn something is to trace its shape or to make a portrait of it.)

¹⁶ If the logic is classical, then, since most (if not all) speakers have inconsistent beliefs, most will have beliefs that imply all existence claims (as well as their negations). So

Quine emphasizes the triviality of his criterion: what exists according to a theory is what it logically implies exists.¹⁷ But, however trivial the criterion may or may not be, applying it certainly need not be trivial, even by Quine's lights. First, it might not be obvious what one believes. It has been claimed, for instance, that various putative beliefs are in fact a form of pretense, a way of making as if one believes. We can seem to hold beliefs that would commit us to the existence of numbers, for instance, but perhaps we are only engaging in a useful make-believe (cf. Field 1980). Second, even when it's clear that one's attitude is belief, it may be unclear what one believes. To take a familiar example, one's belief that the average family has 2.5 children needn't commit one to the existence of the average family (cf. Higginbotham 1985). Whether some belief commits one to the existence of k 's isn't always obvious from the surface form of its expression in natural language. Quine holds that we must first settle on how to represent the sentence in some canonical notation for which what counts as an existence claim and what counts as an instance of logical implication is sufficiently clear.

Applying Quine's criterion of ontological commitment is thus a three-step process. One determines what is believed, one represents it in canonical notation, and then one sees what existence claims are logically implied thereby. This conception suggests two distinct roles an account of semantic competence might be taken to have in reaching ontological conclusions: it might play a role in determining what someone believes (what is, if you will, on the list of beliefs) by establishing or yielding belief in the theorems of the truth-theory itself, or it might play a role in establishing how what

perhaps we should understand the beliefs to be restricted to those that persist through some process of epistemic evaluation.

someone believes (what's on the list) should be represented for the purposes of determining ontological commitment. Let's dub these the belief strategy and the canonical notation strategy, respectively. Presently, in examining these strategies in more detail, we will distinguish various ways an account of semantic competence might establish or yield belief in a truth-theory or might provide guidance about canonical notation. But let's first note how these strategies differ in broad outline. I present both as strategies for generating positive ontological conclusions before turning to the possibility of generating negative ontological conclusions from them.

According to the belief strategy, we can argue:

A believes that S is true, or is at least committed to the truth of S.

An account of semantic competence yields that A believes that S is true iff
P.

So, A is committed to the claim that P.

Canonically representing the claim that P for ontological purposes reveals
that it implies the existence of k's.

So, A is committed to the existence of k's.

Note that the account of semantic competence plays no role here in specifying either what the canonical notation should be or how particular claims should be represented therein.

Its job is to get a certain belief—the belief in a T-sentence—on the table.

According to the canonical notation strategy, we can argue:

¹⁷ Quine 1977, p. 240. The criterion might be in a sense trivial while yet being controversial. Why, for instance, limit the commitment to what's logically implied? And

A believes that S.

An account of semantic competence reveals that the content of this belief,
canonically represented for ontological purposes, is that P.

That P implies the existence of k's.

So, A is committed to the existence of k's.

Note that the account of semantic competence plays no role here in helping specify what we believe: it places no new belief on the table. Its job rather is to supply ontologically relevant information about some belief already on the table.¹⁸

Consider, for example, the alleged ontological commitment to events engendered by the standard treatment of adverbs. The difference in argumentative strategy would amount to whether the account of semantic competence establishes ontologically committing semantic beliefs about events or rather reveals an already existing commitment to events in some non-semantic beliefs expressed using adverbs. Modulo, again, details and variations to be discussed presently, the respective arguments might run in outline:

Belief strategy

A believes that Emma is walking slowly, which commits her to the truth
of [Emma is walking slowly].

of course disputes abound as to what constitutes logic.

¹⁸ Let me underscore that in borrowing Quine's talk of a representation in canonical notation, I'm not assuming his particular views of what the canonical notation should be

An account of semantic competence yields that A believes that:

[Emma is walking slowly] is true iff there's an event that's done by Emma, is a walking, and is slow.

So, A is committed to there being an event that's done by Emma, is a walking, and is slow.

Suppose this statement of A's commitment provides, or is relevantly similar to, its canonical representation for ontological purposes.

There being an event that's done by Emma, is a walking, and is slow implies the existence of events.

So, A is committed to the existence of events.

Canonical notation strategy

A believes that Emma is walking slowly.

An account of semantic competence reveals that the content of this belief, canonically represented for ontological purposes, is that there's an event that's done by Emma, is a walking, and is slow.

That there's an event that's done by Emma, is a walking, and is slow implies the existence of events.

So, A is committed to the existence of events.

Both strategies draw ontological commitments from speakers' beliefs. But they differ both in the role the account of semantic competence plays in generating these

or how it relates to natural language. The phrase, as I'm using it, is at this point just a placeholder for a possible role cognized truth-theories might have in ontological debates.

commitments and in the epistemic status they require of the truth-theory. With the belief strategy, the role of the account of semantic competence is to reveal a semantic belief (that is, a belief in a T-sentence) that, combined with a non-semantic mundane belief (such as the belief that Emma is walking slowly), yields an ontological commitment. This requires, obviously, that the truth-theory be believed. As far as this argument goes, if there were no belief in the truth-theory, no ontological commitment would be forthcoming, even if the truth-theory were itself true. In addition, the strategy requires that the semantic belief and the mundane belief be believed by the same subject. Nothing so far said, however, requires that either belief be true—although, of course, if the person running the argument is herself the believer, then she is indeed committed to the belief's truth.¹⁹

With the canonical notation strategy, on the other hand, the role of the account of semantic competence, as far as the argument goes, isn't to generate a further, semantic belief in addition to a mundane belief. Rather, its role is to reveal something ontologically relevant about what is already believed (that is, about the mundane belief). The canonical notation strategy, unlike the belief strategy, is thus committed to a philosophical claim concerning the particular relevance of a truth-theory to ontological concerns: it presumes that truth-theories provide ontologically relevant information about sentences used to express mundane beliefs. No such premise is deployed by the belief strategy: it just draws out the logical implications of purported beliefs without making any claim in advance for the particular ontological relevance of truth-theories. The canonical notation strategy is thus more committed than the belief strategy in this respect. In another respect, however,

¹⁹ In addition, grounds for thinking the ascribed truth-theory must be true are of course also grounds for believing it, a fact exploited below when we consider the belief strategy

it is less committed. For the canonical notation strategy doesn't require belief in the truth-theory. It suffices, on this view, that the truth-theory does provide ontologically relevant information, whether it's believed or not. Conscious belief in the truth-theory would merely enable one to realize that there is—and has been—such ontological commitment. The person doing the realizing, though, needn't be herself the person ontologically committed. As to whether the canonical notation strategy, requires that the truth-theory be itself true: nothing in the strategy so far requires this, but it's natural to expect that the reasons offered for holding that ascribed truth-theories supply ontologically relevant information will also require that the theorems be true.

Two further points about these strategies. First, conclusions about ontological commitment are not conclusions about ontology: it's one thing for someone to be committed to the existence of k's, and something else (not entailed by the first) for k's to exist. However, someone who acknowledges the commitment in her own case thereby endorses an ontological claim. It's thus important to note that we ourselves, assuming we hold the relevant beliefs, are among the speakers at issue. According to these arguments, given that I believe (for example) that Emma is walking slowly, I am committed to the existence of events. But to endorse such a claim about my commitments is to endorse a claim about what exists.²⁰ Thus will an ontologizing semanticist and whoever accepts her arguments assert, not just that certain speakers are committed to the existence of k's, but also, since they are themselves such speakers, that k's exist.

with the theorist as believer.

²⁰ This is just a variant on Moore's paradox. There is something wrong in saying "I am committed to the claim that P, but it's not the case that P"—unless as a prelude to reconsidering one's commitments.

Second, it's important in applying either strategy that one avoid the natural temptation to read ontological categorizations into truth-theories when they are not there. I have spoken of truth-theories that yield theorems such as: [Emma is walking slowly] is true iff there's an event that's done by Emma, is a walking, and is slow; and I have outlined two strategies for using an ascription of such a truth-theory in order to assign ontological commitment to events. Suppose, however, that the ascribed truth-theory rather yields: [Emma is walking slowly] is true iff there's something that's done by Emma, is a walking, and is slow. Then neither strategy yields ontological commitment to events (absent ancillary premises such as that A is committed to walkings being events), even if A is thereby committed to the existence of what in fact are events. Whether an accurately ascribed truth-theory marks these "somethings" as events, as my exposition has assumed, is an empirical question. However, my discussion doesn't turn on the point: even if the somethings are not so marked, and the speaker at issue is not committed to relevant ancillary premises, one still wants to ask whether an account of semantic competence can yield ontological conclusions regarding, for example, walkings. That said, I will stick to the expositoryly convenient assumption that the ascribed truth-theories advert to such ontological categories.

III. Negative Strategies

These strategies provide ways of arguing for positive ontological conclusions. But what about drawing negative ontological conclusions from an account of semantic competence, conclusions that there don't exist certain entities? One tack would be to base

such conclusions on the absence of successful applications of the above strategies for yielding positive results.

The negative belief strategy would run roughly as follows. The truth-theory ascribed to account for semantic competence doesn't yield a truth-theorem whose right-hand-side (canonically represented) logically implies the existence of k's. So, it's not the case that detaching a truth-theorem's right-hand-side on the basis of belief in its left-hand-side yields a belief that logically implies the existence of k's.

Note, however, that it hardly follows, from the account of semantic competence's not requiring in this manner that k's exist, that k's don't exist. It's not even clear that it provides prima facie reason for thinking they don't. For the fact that there's no application of the belief strategy that yields an ontological commitment to k's doesn't mean that there couldn't be ontological commitment for some other reason. It's not part of the belief strategy to claim that arguments of this form are the only way one can establish ontological commitment; nor would such a claim be justified.

Further, it's unlikely that no truth-theorem's right-hand-side logically implies the existence of k's. What of the truth-theorem for 'There are k's'? A negative belief strategist can point out that such a truth-theorem only enables a positive application of the belief strategy if this sentence expresses something already believed, which it doesn't in her case. But this would only show that the account of semantic competence provides no grounds to countenance k's to those who don't already countenance them. It wouldn't show that k's don't exist.

Might a negative canonical notation strategy take one further? Suppose that no sentence is such that its canonical notation, as revealed by an account of semantic

competence, logically implies quantification over k's. Then (the argument would run) no matter what one believes, no sentence expressing that belief is ontologically committing to k's. So, on this supposition, it would seem impossible to be ontologically committed to k's. A proponent of the canonical notation strategy might thus argue that k's don't exist (and not merely that the account of semantic competence provides no reason to think they do exist). It would seem, however, again too much to expect that no sentence, when represented in canonical notation, logically implies the existence of k's. What of the truth-theorem for 'There are k's'? Note, further, that if no sentence, when canonically represented, logically implied the existence of k's, one would seem unable even to express the position that k's exist. Surely, Ludlow (for instance) wouldn't deny that an opponent who holds that past events do exist can express this in English.²¹ If so, then there must be a way to express the position, one whose representation in canonical notation logically implies the existence of past events. A natural candidate would seem to be the sentence 'There exist past events.'²²

So, suppose S is a sentence in the speaker's idiolect whose canonical representation logically implies the existence of k's. Still, one might deploy the negative canonical notation strategy to draw a negative ontological conclusion by reasoning as follows. The account of semantic competence shows that practically no sentence (whether one expressing a claim in metaphysics, or in physics, or what have you) is such that it's ontologically committing to k's. So, if a sentence that is ontologically committing to k's (such as, let's suppose, 'There are k's') expresses something a speaker

²¹ I am indebted here to a remark made by Zoltán Gendler Szabó.

²² If the semantics of such "existential-there" constructions are fit into the framework of Generalized Quantifier theory, however, the question arises whether their truth-theorems'

believes, then it expresses a belief without warrant: for no matter what else she believes, her other beliefs ex hypothesi fail to logically imply it.

This argument, however, assumes that a belief is without warrant unless it is logically implied by other beliefs. This is not in general the case: other beliefs may lend warrant to a belief without logically implying it, and (at least on many views) it is possible for a belief to be warranted without its warrant resting on that of any other belief. So, to reach the conclusion that a belief in k's would be without warrant, one would need to supply—in the given case—reason to think that warrant of these other sorts is not to be had.

There is thus a difference in the power of the positive and negative strategies. Instances of the positive strategies, supposing they're sound, indeed establish ontological commitments. Instances of the negative strategies, however, even supposing they're sound, can at best lend support to ontological denials.

But are the belief and canonical notation strategies (and thus their negative counterparts) sound? Before we address this question in earnest in the next two sections, let me mention—to set aside—two ways one might prevent an instance of any of these arguments from generating its ontological upshot. First, one might abandon the mundane belief that begins the argument. For instance, one might abandon the belief that Emma is walking slowly—perhaps indeed in light of its alleged ontological commitment to events. One might then need to explain away the appearance to the contrary that you do, or did, maintain such a belief—perhaps by adverting to some ersatz attitude. Second, one might abandon the claim that the t-sentence adverted to is a theorem of the truth-theory

right-hand-sides in fact logically imply the existence of k's or, rather, only do so with the help of some set theory.

cognition of which in fact explains semantic competence. For example, a disbelief in events could provide grounds for proposing an alternative treatment of adverbs (supposing one could be found that accounted for the relevant data).

The availability of these moves—along with both strategies’ inability to establish negative conclusions—underscores that the study of semantic competence can’t play the role of first philosophy with respect to ontological questions, providing the sole discipline within which such metaphysical disputes must be (empirically) adjudicated. As Ludlow puts it, on his approach to metaphysics

. . . it becomes a question of simultaneously solving the constraints of semantical theory and constraints imposed by our metaphysical intuitions. Neither set of constraints is a priori privileged; the resulting theory must satisfy both.²³

Still, it is the view of the ontologizing semanticist that, given the mundane belief and given the truth-theory cognized, positive ontological commitments follow. Otherwise put, it’s not the case (on this view) that, given the mundane belief, the truth-theory yields merely prima facie evidence of positive ontological commitment, evidence that can be outweighed without calling into question the attribution of the truth-theory itself. When it comes to applications of positive instances of the strategies, “simultaneously solving the

²³ Ludlow 1999, p. 70. Perhaps the constraints go beyond those imposed by semantic theory and by our “metaphysical intuitions.” Compatibility with our best physics, for example, would seem relevant to questions concerning time. (See Ludlow 1999, pp. 2-3, for remarks on the bearing of physics on the metaphysics of time. Also, Ludlow 1999, p. 70, mentions the possible bearing of physical theory on the question of the existence of properties.)

constraints” requires abandoning the mundane belief, altering the truth-theory, or accepting ontological commitment.²⁴

At least so it is claimed. We now examine more closely the belief strategy and the canonical notation strategy in turn. For a positive ontological commitment in fact to follow, given some mundane belief and a truth-theory, an ontologizing semanticist (I argue) would have to defend some further claims—but differing claims, depending on which strategy she employed. The points at which one might resist these strategies correspondingly differ.

IV. The Belief Strategy

There are two ways one might run the belief strategy, depending upon who the believer is—the speaker of the language or the semantic theorist.

A. Speaker’s Belief Strategy

Recall our sketch of the belief strategy:

A believes that S is true, or is at least committed to the truth of S.

An account of semantic competence yields that A believes that S is true iff

P.

So, A is committed to the claim that P.

²⁴ The situation is also holistic in that a failure to get the theory to mesh with metaphysical intuitions (or whatever else) could in principle provide a reason, not just to abandon a particular semantic proposal, but to abandon the larger cognitivist truth-theoretic approach—an option, however, that we put to one aside above.

Canonically representing the claim that P for ontological purposes reveals
that it implies the existence of k's.

So, A is committed to the existence of k's.

When the believer is the speaker of the language, 'yields' in the second premise means: provides sufficient grounds to conclude. (As we'll see, it must be understood differently when the strategy is applied to a theorist *qua* theorist.)

The speaker's belief strategy makes three assumptions that are especially worth noting, since all three might be resisted. First, it assumes that the canonical representation of the truth-theorem's right-hand-side does indeed reveal it to be ontologically committing. As is familiar, this is a quite substantial assumption. Ludlow, for instance, defends it against the objection that the meta-linguistic quantification might be construed substitutionally.²⁵ In addition, there is the concern, with negative applications of the strategy, that the assignments of truth-conditions heralded as avoiding ontological commitment might actually reveal ontological commitment upon canonical representation. This worry might be raised, for instance, about the assignments of truth-conditions Ludlow suggests in order to avoid quantification over past and future events in handling tense. Consider the t-theorem: [Emma was walking] is true iff the ILF of [Emma is walking] was true. Belief in the left-hand-side, combined with belief in the t-theorem, commits one after all to past events, if the canonical representation of the right-hand-side is (say): there exists a time prior to the present at which the ILF of [Emma is walking] is (timelessly) true. If this were the case, Ludlow would be blocked from

²⁵ Ludlow 1999, p. 74-6, follows Parsons 1971 in suggesting that such quantification doesn't avoid ontological commitment.

pursuing a negative speaker's belief strategy to support denial of past and future events. Ludlow, however, suggests that tense morphemes be construed as evidential-modals (roughly, the claim that an ILF was true would be a claim about the evidence we presently possess for it).²⁶ If this idea were to pan out, perhaps it could assuage this concern. In any event, on the belief strategy it's not claimed that what underwrites the first assumption is the account of semantic competence. For this reason—and because worries about how to interpret the quantifiers are familiar—I won't discuss it here further.

The speaker's belief strategy's second assumption is that cognition of a truth-theory involves belief on the part of the speaker. In resisting it, one might draw attention to the fact that the attributed states tend to exhibit many of the marks of modularity.

- (1) They are in principle inaccessible to consciousness. Even if one takes a course in which one learns a correct semantic theory for one's own language, one does not thereby raise to consciousness these states. Rather, one comes to know of these states from a third-person perspective, much as one would if one studied how the human digestive system works.²⁷

²⁶ Ludlow 1999, pp. 156-63. Ludlow is mainly concerned to suggest that we might dispense with the grammatical category of tense. I am in effect adapting his remarks into a suggestion about the lexical entries for tense morphemes. (Ludlow also speaks of the simple tense morphemes and 'when' as each being "a kind of primitive" (p. 112). I believe he's here highlighting that the right-hand-sides of their clauses are with respect to these expressions homophonic, as opposed to the standard treatment's clauses, which introduce quantification over events.)

²⁷ It might be objected that not all of the truth-theory is in principle inaccessible to consciousness: perhaps the relation speakers bear to much of the truth-theory isn't one of belief, but the relation they bear to the relevant truth-theorems is, which is all the speaker's belief strategy would require. The right-hand-sides of truth-theorems, however, are typically far from homophonic: as we've seen, for example, Ludlow's treatment of tense morphemes treats them as predicates of ILF's. To consciously work through the derivation of such truth-theorems is not to bring an existing state to consciousness—even

- (2) They are informationally encapsulated. Much semantic information seems proprietary to the language module. Certainly, given (1), it plays no conscious role in reasoning.
- (3) Relatedly, they needn't be "conceptualized"—that is, the speaker needn't grasp the concepts deployed in these sophisticated truth-theories.
- (4) Semantic comprehension is to a large extent automatic—almost a reflex. To use Hume's (1779, III) example, we can't but hear significance in a rush of the wind, if it sounds sufficiently like phonemes of a language we speak.

Let's lay out in stages how these facts might be deployed to resist the assumption that speakers believe the truth-theory (that is, believe its theorems, on the basis of believing its axioms and applying its production rules).

First, given these marks of modularity, the resister might question attribution of belief to the speaker. Ascribing cognition, it will be suggested, involves an attribution to some sub-personal cognitive system of the speaker rather than an attribution of belief—or any other attitude—to the speaker herself. The ontologizing semanticist, however, can reply that exhibiting the marks of modularity doesn't preclude the ascription of an attitude to the person: a person can have an attitude in virtue of some sub-personal cognitive system of hers having one.

Second, the resister might suggest that these marks of modularity, even if they don't preclude attribution to the speaker, indicate that the attitude ascribed to the speaker falls short of belief. The ontologizing semanticist here can reply that truth-theories,

if the resulting state is identical in content to one causally relevant to one's own linguistic behavior. This is not to deny, however, that conscious states can significantly affect aspects of the cognized truth-theory—e.g., when an axiom for some expression is altered

insofar as they are believed, are believed largely tacitly, so that the marks of modularity merely show what kind of belief is involved.

Third, the resister might suggest that the attitude—however classified—is not one relevant to ontological commitment, so that even if it is a kind of belief, it's not a kind that's ontologically committing. At this point we encounter an unclarity in the Quinean conception of ontological commitment.²⁸ Do the beliefs relevant to ontological commitment include tacit beliefs? If the ontologizing semanticist is to shoulder the burden of proof here, she must provide reason to think they do. One prima facie obstacle she faces is that we don't generally epistemically assess an agent on the basis of states in principle inaccessible to consciousness—for example, a linguist who on the basis of the latest state-of-the-art theorizing self-ascribes a syntactic theory incompatible with the one she in fact cognizes isn't on that account irrational.

Further, the resister might support her denial that cognition involves ontologically committing belief by offering an alternative construal of the attitude of cognition. Let me offer an example, which may or may not have empirical legs: the point is to show that it's not trivial to claim that cognition involves ontologically committing belief. So, the resister might propose that cognizing be understood as a kind of tacit pragmatic presupposition. Someone pragmatically presupposes P, roughly, iff she takes P for granted for present conversational purposes and takes others party to the particular conversation likewise to take P for granted (cf. Stalnaker 1974). Note that pragmatically presupposing P—even explicitly—is rationally compatible with believing that not-P. P

in response to one's realization that one's understanding of the term significantly deviates from that of others.

can be something that is just being accepted for the nonce, perhaps a supposition, something the parties are willing to make as if they believe. The suggestion, then, is that speakers—at least as a default—tacitly expect others with a sufficiently similar idiolect to grant as part of the conversational background that the expressions they use have the semantic values assigned by their truth-theories.

Now, since one can without inconsistency pragmatically presuppose the existence of what one believes does not exist, pragmatic presupposition is not ontologically committing. So, if cognition is tacit pragmatic presupposition, then cognition is not ontologically committing. Though we might, we needn't add that cognition is not a kind of belief: perhaps it's a kind of tacit belief (maybe even this is what tacit belief is generally, though we needn't go that far either). But then it would follow that tacit belief—and thus belief generally—needn't be ontologically committing. In sum, if cognition is a kind of pragmatic presupposition, then, if it's a kind of belief at all, it's not a kind of ontologically committing belief. The ontologizing semanticist pursuing the speaker's belief strategy would therefore need to show that the explanation of linguistic behavior requires ascriptions of attitudes that go beyond cognition construed in terms of pragmatic presupposition.

It's not obvious, however, how this might be done. If we consider explicit attitudes, then what distinguishes belief from as if belief are their associated epistemic and practical responsibilities. Someone who explicitly believes something must be responsive to reasonable challenges to the claim believed; someone who explicitly makes as if she believes something is responsible only for the reasonableness of so making, not

²⁸ An unclarity, that is, for cognitivists who would adopt a Quinean conception of ontological commitment: Quine himself didn't countenance such attitudes, so didn't need

for the claim's warrant. This distinction only applies, however, to attitudes for which the agent is responsible in the first place. Attitudes in principle inaccessible to consciousness, however, are not under an agent's proper control, and so an agent is thus not typically responsible for them. Again, explicit belief and explicit as if belief differ in their practical upshot. The difference is not, as it's tempting to think, that as if belief (for example, supposition), unlike belief, is disengaged from action, but rather that as if belief requires more than belief does in order to provide reason for action. Supposition can provide reason for action, for example, when one also believes that acting on the supposition will tend towards success, as can happen in cases of reasonable instrumentalism.²⁹ Now, the tacit attitudes we are interested in are ascribed in order to explain behavior, so it's of course assumed that they are not disengaged from action. One might then think to distinguish tacit pragmatic presupposition and some stronger, more belief-like cognition on the basis of whether there are grounds for ascribing as well the tacit belief that acting on such a presupposition will tend towards success. But in any instance where one might consider ascribing pragmatic presupposition inaccessible to consciousness, the very fact that the ascribed attitude would be implicated in behavior—the fact, that is, that one's not just ascribing the attitude, but also hypothesizing for it a certain role in the speaker's cognitive economy—would provide grounds for claiming that the attitude is tacitly held to be a basis for action. Thus, both when we advert to epistemic and to practical considerations, we seem unable to exploit at the tacit level what distinguishes belief and as if belief at the explicit level. It's thus unclear how the ontologizing semanticist might

to consider whether they generate ontological commitments. See Quine 1970.

²⁹ One can have reason to act on a supposition for other reasons as well—for example, as a way of discovering whether the supposition is true, something common in scientific and other inquiries.

establish that, contrary to our suggestion, explaining linguistic behavior requires the positing of ontologically committing attitudes—a burden they must bear so long as less committing alternatives seem available.³⁰

I turn now to the third important assumption the speaker's belief strategy makes: that speakers believe, not only (for instance) that Emma is walking slowly, but also that [Emma is walking slowly] is true—for it is only the latter belief that allows us to infer the right-hand-side of the relevant biconditional. But this assumption can be resisted for reasons parallel to those given above. A speaker of the relevant language who believes that Emma is walking slowly and who possesses the concept of truth³¹ is committed to the claim, and typically will believe, that 'Emma is walking slowly' is true. But arguably most speakers lack the syntactic concepts necessary even to form explicit de dicto beliefs involving [Emma is walking slowly], even if they may cognize such claims as that [Emma is walking slowly] is true. Of course, this isn't true of all speakers: theorists do possess the necessary concepts. The third assumption is thus not unreasonable when attention is restricted to theorists.³² It's then natural to consider restricting the second assumption—concerning belief in the truth-theory—to theorists as well. This leads to the second version of the belief strategy.

³⁰ One might think that an appeal to “making as if” or pretense can be challenged on empirical grounds if, prior to their developing the ability to engage in pretense, children exhibit the behavior the pretense is invoked to help explain. Suppose for example children prove adept as using adverbs prior to gaining the ability to pass the false belief test. This objection, however, would not seem relevant to the positing of tacit (or subpersonal, or modular) pretense, which could antedate a person's being able to engage explicitly in make-believe.

³¹ Soames 1989, p. 578, argues that knowledge of truth-conditions is not necessary for semantic competence, since “young children and unsophisticated adults can understand lots of sentences without understanding ‘true’, or any corresponding predicate.”

B. Theorist's Belief Strategy

Suggesting that cognition does not entail belief—at least in any sense that ontologically commits—sometimes inclines ontologizing semanticists (I have found) towards the following response: “But I do believe the biconditional. I think what it says is true!” This response can be understood in two ways. First, she might mean simply to endorse the biconditional, and thus be ontologically committed to k's, without suggesting that her acceptance of the biconditional is based on an account of semantic competence. (Perhaps it's rather something she believes on independent grounds.) This is obviously not of interest to us, since we're exploring the move from accounts of semantic competence to ontology. Challenging this move doesn't by itself preclude the possibility of accepting even on the basis of this biconditional that k's (events, say) exist. The question is whether the biconditional's place in an account of semantic competence is what imbues it with ontological significance.

The second way of construing the response is of interest to us. The ontologizing semanticist might be pointing out that the account of semantic competence doesn't just claim that speakers qua speakers (tacitly) believe these biconditionals: it asserts such claims itself, and so the biconditionals are among what theorists believe qua theorists. But linguistics is a science, whose deliverances should be taken seriously. So, if theorists believe such things on the basis of their accounts of semantic competence, then those beliefs are as ontologically committing as the deliverances of geology.

³² An alternative reply might be that the speaker can cognize, even if she doesn't believe, that 'S' is true iff [S] is true and so, from the truth-theorem for [S], derive one free of syntactic concepts and thus eligible for belief.

This provides the second way of running the belief strategy, now focused on the theorist as believer instead of on the speaker of the language. To run the strategy this way we must reconstrue the argument's second premise—that an account of semantic competence yields that A believes that S is true iff P. Before, when A was some speaker, 'yields' meant: provides sufficient grounds to conclude. But when A is a theorist *qua* theorist, its import must be: has as a result. For, according to the theorist's belief strategy, an account of semantic competence provides sufficient grounds for theorists to conclude that [S] is true iff P, and therefore the theorist believes that [S] is true iff P.³³ Thus adjusted and assuming it sound, the belief strategy can only generate ontological commitments for the theorist, since the subject of the semantic belief (belief in the T-sentence) must be the same as the subject of the mundane belief used to reach the T-sentence's right-hand-side. The ontological commitment can be extended, however, to non-experts who accept the testimony of theorists.

Now, one basis for resisting the theorist's belief strategy carries over from our discussion of the speaker's belief strategy: viz., the concern that the t-theorems' right-hand-sides, when canonically represented, might not have the ontological consequences claimed for them. Perhaps less obviously, one might also retain the concern that ascription of cognition of a truth-theory does not establish, or provide grounds for, belief in it—albeit now belief among theorists.

³³ The role of the account of semantic competence, on the speaker's belief strategy, is to justify ascription of a truth-theory to a speaker; its role on the theorist's belief strategy is to justify the truth-theory itself. A theorist can indeed ascribe a truth-theory to herself qua speaker (third-personally, if you will), as running the speaker's belief strategy concerning oneself would require. But, according to the theorist's belief strategy, her theorizing also provides grounds for endorsing what is thereby ascribed—and thus provides grounds, not just for ascribing belief to herself, but for belief.

What are the deliverances of accounts of semantic competence? Do they include such claims as:

[Emma is walking slowly] is true iff there is an event e such that Emma walks in e and e is slow,

as the argument above would have it, or are they rather limited to such claims as:

speakers of the language cognize that [Emma is walking slowly] is true iff there is an event e such that Emma walks in e and e is slow?

That is, does the account deliver the theorems of the truth-theories it ascribes, or does it just deliver the ascriptions? If only the latter, then the theorist's theoretical belief does not commit her to the existence of events: of course in general one needn't believe what one ascribes to others; but further, if what one's ascribing is cognition, then even self-ascription doesn't require that the theorist also believe what is ascribed. The ontologizing semanticist pursuing this strategy must therefore maintain that her theorizing does indeed deliver claims of the first sort and not just of the second—so that such claims are at least among what she herself believes *qua* theorist. She is committed, that is, to holding, not only that the truth-theory's theorems are themselves true, but further that they are among what theorizing about semantic competence establishes.

No doubt many theorists would embrace this commitment. I want to suggest, however, that it's unclear why a truth-theory's role in a cognitivist account of semantic

competence should require that the theory (that is, its theorems) be true, as opposed to its being a useful way for speakers sub-personally to model how things—or certain things—are for certain cognitive purposes. Let me run through some of the reasons one might give in favor of the truth of ascribed truth-theories. I'm not sure whether any succeed. But even if one does, it's important to note that these various lines employ varying—in some cases conflicting—premises. It would be important to know to which premises an ontologizing semanticist must commit herself.

1) It might be thought a theorem's failure to be true might undermine explanations of linguistic behavior that advert to it. But if ascribing cognition of a truth-theory suffices to explain the relevant aspects of linguistic behavior, then it would seem to go beyond what explanation requires to claim further, on the basis of its role in this explanatory hypothesis, that the ascribed truth-theory be true. Compare a more mundane example of action explanation adverting to cognitive states: It would seem that what explains my walking into the kitchen is my belief that there's food there and my desire for food; whether my belief is true is irrelevant to the accuracy of the explanation. If this holds as well for explanations of linguistic behavior, then not only—as was suggested above—can the speaker (or, if you will, the speaker's semantic module), compatibly with an accurate account of semantic competence that adverts to k's, remain agnostic about their existence, but so can the theorist (the explainer).

2) It might be suggested that, although the explanation of my walking into the kitchen doesn't require the truth of my belief, an explanation of its success as a means towards satisfying my desire for food does. Similarly, if we want to account for the success of our linguistic behavior, we must ascribe truth to the relevant cognitions:

accounts of semantic competence yield ontological conclusions because of this competence's role in successful action. We must make explicit therefore an additional premise in the ontological argument to the effect that language use (at least of the bit of language at issue) leads, or tends to lead, to the satisfaction of the end(s) for which it's undertaken—or at least is such that, insofar as it doesn't, the blame lies not with the relevant T-theorem's lack of truth.

But it's not the case generally that an action's (non-accidental) success requires the truth of the beliefs or belief-like attitudes that in part explain the action. If it were, it would preclude any case of successful instrumentalism—any case where P is false but a non-believing (perhaps agnostic) agent's making as if P and acting accordingly tends to satisfy her desires, given what else she believes and supposes, and thus tends to render her so acting successful.³⁴

Perhaps, though it's not generally the case that an action's success requires the truth of the attitudes that in part explain it, this is required for successful linguistic behavior, or at least that range of it involving the allegedly ontologically salient constructions. It's unclear, however, how an ontologizing semanticist might support this. Evolutionary considerations don't generally favor true beliefs (or, cognitions) over instrumentally reasonable ones; so, why should they here? And Davidsonian considerations of rationality, charity, and/or humanity—assuming they are granted—likewise don't obviously preclude a range of reasonable, but as it happens false,

³⁴ The suggestion is that the speaker be treated as an agnostic instrumentalist about the truth-theorems, not about mundane beliefs such as that Emma is walking slowly. And this is consistent with the theorist being a realist about both ascriptions: the theorist can say that the speaker really is (in a causally relevant way) tacitly supposing the truth-theory and believing the mundane belief. It's not being suggested that the theorist be (or, be understood as) an instrumentalist.

presuppositions within what we tacitly cognize. Even if we suppose that the speaker does believe the truth-theory ascribed by a cognitivist account of semantic competence, it's not obvious that such Davidsonian considerations dictate that these beliefs be among those others must believe as well in order to interpret the speaker.³⁵ In any event, it seems a substantial claim, in need of defense, that explaining successful linguistic behavior requires the truth of a truth-theory's theorems.³⁶

3) One might try arguing on individualist grounds that the cognized truth-theory must be true. We need to distinguish here two kinds of individualism: linguistic individualism and psychological individualism. Linguistic individualism, as applied to the kind of theorizing at issue, is the thesis that truth-theories ascribed to account for semantic competence should aim to capture the truth-conditions individuals' semantic modules assign to sentences, and not the truth-conditions sentences possess in some supposed "public" language such as English. Psychological individualism is the thesis

³⁵ Davidson 1973 claims that interpreting the speech of another requires taking "most" of her beliefs to be true, or at least requires that beliefs must be ascribed to her so as to render her rational or sufficiently human-like. Davidson 1983 argues further that "most" of beliefs must in fact be true. (Such considerations, if sound, might provide a way of moving from claims concerning ontological commitment to claims concerning ontology.)

³⁶ It might be suggested that, even if explaining the success of linguistic behavior doesn't require the truth of the ascribed truth-theory, still a truth-theory's role in such explanations provides prima facie reason for believing the theory true. Even this, however, is not clear. It's plausible that the success of our scientific endeavors does provide at least prima facie grounds for the truth of the relevant beliefs. This is because success here includes coming up with beliefs that yield accurate predictions. But linguistic behavior is arguably different. How best to characterize linguistic success is not obvious. But suppose, following a familiar line, we say that at least in a range of core cases it involves one's audience understanding one's speech act. It may suffice to explain such success as there is to note that the language-users cognize sufficiently similar truth-theories and are otherwise relevantly similar in their linguistic (and more broadly their cognitive) competence. But it's not obvious why this mutuality should provide prima facie grounds for the truth of the cognized truth-theory.

that what psychological states one's in is determined by one's internal, physical states.³⁷ Applied to the case at issue, it implies that the content of the cognized truth-theory is individuated individualistically. Both theses are controversial: some ontologizing semanticists would not want to rest their claims upon them.³⁸ But even if they are accepted, does either support the truth of the ascribed truth-theory?

If linguistic individualism is true, then it can seem that there's no room for a gap between what you cognize and its truth: that you cognize it makes it true. But this is not the case, at least as I've characterized linguistic individualism. I characterized linguistic individualism as the thesis that truth-theories ascribed to account for semantic competence should aim to capture the truth-conditions individuals' semantic modules assign to sentences. Nothing in this characterization suggests that, if you get right what truth-conditions the module assigns, then those truth-theorems must be true. It's consistent with linguistic individualism so characterized that the semantics module assign false truth-conditions. Perhaps, then, linguistic individualism shouldn't be so characterized. Perhaps the relevant thesis is rather that truth-theories ascribed to account for semantic competence should aim to capture the truth-conditions sentences in fact have in a speaker's idiolect.³⁹ The question then becomes why we should accept this thesis, especially since my weaker formulation suffices to distinguish the linguistic individualist's project from that of the "public" language semanticist. It allows, for

³⁷ I take this formulation of psychological individualism from Larson and Segal 1995, p. 530. As noted below, I depart from their characterization of linguistic individualism.

³⁸ While Larson and Segal 1995, pp. 525-42, accept both forms of individualism, Ludlow 1999, pp. 152-6, rejects the second.

³⁹ This is closer to Larson and Segal's characterization of linguistic individualism: "the semantic features of a person's language are whatever his or her internalized semantic theory states them to be." (1995, p. 529)

instance, that the truth-theory correctly ascribed to some speaker may relevantly depart from that correctly ascribed to some expert.

A truth-theory might fail to accord with what a sentence means in more than one way. The linguistic individualist is anxious that a failure to align with what expressions mean in a public language (as when I take ‘arthritis’ to refer to inflammation either in joints or in muscles) not be a reason to reject a proposed hypothesis concerning the truth-conditions assigned by a speaker’s semantic module. But one can consistently reject the relevance of that sort of misalignment while allowing that some of the truth-theorems correctly ascribed are in fact false and in that sense don’t accord with what the expressions mean—for example, because they assert the existence of stuff supposed reasonably, but perhaps falsely, to exist.

Consider now psychological individualism—the thesis that what psychological states one’s in is determined by one’s internal, physical states. Suppose psychological individualism were upheld and, further, that what truth-conditions your sentences in fact possess is determined by your psychological states. Perhaps one might even consider sentences in your idiolect having the truth-conditions they do to be a psychological state of yours. Then it might be thought required that the truth-conditions assigned by the semantics module be true; for what else could determine what the truth-conditions in fact are besides the truth-theory you cognize? But does it in fact follow that, if the truth-conditions of a sentence in a speaker’s idiolect are determined individualistically, then they are determined only by the truth-theory the speaker cognizes? Not necessarily. Nothing so far said precludes the truth-conditions’ being determined in part by other states. For example, nothing so far said precludes a speaker’s cognizing that [Necessarily

P] is true iff, in all possible worlds, P, while—given the role this cognition plays in the speaker’s larger cognitive economy and the rest of the speaker’s attitudes—[Necessarily P] is true in the speaker’s idiolect iff, supposing there are possible worlds, in all of them, P.

4) One might try to argue on externalist grounds that expressions must have the properties one cognizes them to have—or at least that this is so for certain properties. It’s unclear how this might go, but perhaps, as applied to ontological commitments, the thought would take something like the following form: our mental states have the content they do in part owing to relations to the external world; so, if the world lacks a certain kind of entity, it’s hard to see how they could be posited in our cognitions.⁴⁰ Again, it would be surprising if an ontologizing semanticist had to commit herself to externalism, given the divergence ontologizing semanticists in fact display on the issue; and, if she does, then it must be noted that the doctrine is controversial. But, also again, it’s unclear that externalism really has this upshot. Externalists generally need to say something about empty expressions, including empty kind-terms; but the argument given here would seem to imply that there can’t be such terms. If an externalist ontologizing semanticist running this line wants to avoid this result, then she must rest the argument upon particular features of the putative kind at issue—perhaps on its being an expression for an in some sense fundamental ontological category. It remains unclear to me, however, how this would best proceed.

⁴⁰ This adapts a remark of Ludlow. To a B-theorist who suggests that temporal indexicality is found only in our thoughts, Ludlow replies: “. . . psychological states (particularly perceptual states) are individuated in part by relations to the external world. . . if the world is not tensed then it is difficult to see how our perception of the world could be tensed.” (1999, p. 96) Perhaps Ludlow intends his argument to turn on particular features of tense, however, and would object to its being generalized.

5) It might be suggested that, if truth-theories needn't be true, then we are deprived of the primary basis for ascribing one truth-theory as opposed to another. A truth-theory is assessed in large part in terms of the truth of its theorems: our ability to test the theory's claims in this way is what gives epistemic purchase to our ascriptions. This is just what becomes of Tarski's Criterion T when one transforms his project of defining truth for a language given "translations" of its sentences into the meta-language into the project of specifying the sentences' truth-conditions using the meta-linguistic concept of truth: we want that, for each sentence, truth-conditions get assigned that we can verify as the right truth-conditions. Getting the truth-conditions right may not exhaust the evidence available in assessing ascriptions of truth-theories (that it doesn't, since it matters also whether the ascription jibes with neuroscience, developmental psychology, psycholinguistics, etc., is a crucial claim of cognitivist theorists), but surely it's a core part.

One can deny that correctly ascribed truth-theories need be true, however, without allowing that anything goes. We've been exploring the possibility that cognized truth-theories might fail to be true in a particular way. Cognition might be something like tacit pragmatic presupposition. In particular, it might be that the claims of the truth-theory are accepted under the supposition that certain kinds of entities exist. We, or our semantic modules, perhaps make as if there are k's (e.g., possible worlds) and then, under this supposition, assign truth-conditions to certain sentences (e.g., to modal sentences). Perhaps these suppositions enable us to use and interpret these expressions in a way particularly well-suited given aspects of our cognitive apparatus—for example, by putting them in a form amenable to certain sorts of combinatoric treatment. If our

linguistic behavior and other relevant evidence can be well accounted for given the ascription of such a truth-theory, it needn't matter that the theorems yielded—once the supposition is dropped—might in fact be false (e.g., if it turned out that there are no non-actual possible worlds). It might matter only that it be a successful, or reasonable, strategy for the speaker, or her semantic module, to make as if they're true. As long as we can differentiate candidate truth-theories in terms of the success or reasonableness of speakers making as if they are true, then we needn't limit ourselves only to true theories.

6) It might be claimed that articulating what semantic properties sentences in fact have just is the job—or at least a primary job—of semantic theorizing. Getting sentences' truth-conditions right is at least a part of this; so, a truth-theory whose assignments of truth-conditions are not themselves all true has fallen down on the job.

There are many projects, however, that go by the name 'semantics,' and it's not obvious that a theory that plays a part in one must also play a part in the others. Our focus has been on the project of accounting for semantic competence and thus contributing to an explanation of linguistic behavior. It's claimed that such an account involves the ascription of a truth-theory. I've been exploring whether a truth-theory could fulfill this function even though it's not the case that all of its theorems are true. It's consistent with this to maintain that there are other projects whose success does require the deliverance of true truth-theorems. It's thus a move to go from the legitimacy of these other projects to the requirement that the truth-theory ascribed in accounting for semantic competence be true. Scott Soames 1989, p. 575, holds that it's the job of semantics to articulate "what information is encoded by sentences relative to contexts," regardless of what semantic competence—speakers' mastery of those expressions—consists in or involves. Likewise,

one might pursue an account of semantic competence without worrying whether one's results yield an accurate assignment of truth-conditions or can be converted into statements of meaning or an articulation of encoded information in Soames' sense.

If an ontologizing semanticist claims otherwise, if she maintains that her ontological conclusions are legitimate because a truth-theory ascribed in accounting for semantic competence must also answer to some version of what we might call the meaning-stating project, then she commits herself to several substantial claims. Most obviously, she commits herself to the legitimacy of the meaning-stating project. Note that one may question its scientific bona fides without questioning that of an account (in particular, a cognitivist truth-theoretic account) of semantic competence—perhaps on the ground that it's unclear what explanatory work such a theory is supposed to do, if accounting for linguistic behavior doesn't require it.⁴¹ Of course, merely mentioning the possibility of such a challenge doesn't constitute one. But the point is just that the ontologizing semanticist pursuing this line commits herself to a theoretical goal about which other lines can remain silent.

In addition, the ontologizing semanticist commits herself to the convergence of the meaning-stating project and the project of accounting for semantic competence. More specifically, she must hold, first, that the statement of what sentences mean must take the

⁴¹ Whether an empirical explanation of linguistic behavior even requires a theory deploying the concept of meaning (whether a "folk" concept or one adapted for theoretical purposes) would seem to be itself an empirical question. Kölbel 2001, p. 614, claims that it's a "prima facie absurd view that a theory of meaning for a language does not say what any sentence of that language means." I disagree if a theory of meaning is whatever is ascribed as part of a cognitivist account of semantic competence. Cf. Gross submitted.

form of a truth-theory;⁴² and, second, that the truth-theory that answers to the demands of the semantic project must moreover be the same truth-theory as that deployed in accounting for semantic competence.

One might question these assumptions by pointing out that the truth-theories employed in accounts of semantic competence are constrained in ways not obviously relevant to the task of stating what sentences mean. We saw that developmental data and acquired deficit studies arguably favor Ludlow's treatment of tense over the standard view; but should this matter for stating what information expressions encode? If in the end the evidence supporting Ludlow's alternative was all psycholinguistic, might the meaning-stating theory continue to use the standard view of tense while acknowledging reason to think such an approach inappropriate for accounts of semantic competence? Why generally should psychological data constrain the project of stating what sentences mean? Could there be combinatorial techniques—such as hanging modifiers off of events—that our minds exploit, but that fit less well with, or are just not required by, the meaning-stating project?

Even the requirement of compositionality looks quite different depending on whether it's viewed through the lens of the meaning-stating project or rather the project of accounting for semantic competence. In accounting for semantic competence, it's the constraint of learnability—together with a desire to explain productivity, creativity, and systematicity—that arguably forces the theory to be compositional. But it's not clear that these considerations are relevant to the project of stating what information sentences

⁴² Soames 1989 is an example of someone who would reject this. He maintains that a theory answering to the meaning-stating project should take the form of a compositional theory assigning structured propositions to sentences relative to a context. Truth-

encode. If compositionality is required here, it would seem to be for different reasons: perhaps owing to the limitations of the theorist (who otherwise couldn't finish stating her theory), or perhaps because it is held, on grounds other than those invoked in the context of accounting for semantic competence, that what an expression means—what information it encodes—is what it systematically contributes to expressions of which it is a constituent. Compositionality theses, what's more, come in various strengths; so, even if both theories that play a role in accounting for semantic competence and that state what sentences mean must be in some sense compositional, the differing grounds could motivate different compositionality theses in each case.

7) One might worry that permitting untrue truth-theorems deprives theorists as well of another source of evidence for their ascriptions. Particular truth-theoretic proposals are often supported by appeal to certain valid inferences that they “capture.” Thus, as evidence for an eventish treatment of adverbs, it is often noted that the truth-conditions assigned ‘Emma is walking slowly’ (that there exists an e that is done by Emma, is a walking, and is slow) logically imply those assigned ‘Emma is walking’ (that there exists an e that is done by Emma and is a walking). The worry is that, if the truth-theorems don't give the correct truth-conditions for the target sentences, then it's unclear why the way their right-hand-sides are logically related should be relevant to the inference from one target sentence to the other. But if it isn't relevant, then the truth-theory doesn't capture the inference (at least, not in a sense that matters to what capturing the inference is supposed to explain) and so a significant source of evidence for one truth-theory over another—that it does capture these inferences—seems lost.

conditions are then assigned to such propositions. (Soames argues as well that knowledge of a truth-theory is neither necessary nor sufficient for semantic competence.)

What's meant here by "capturing" a validity? Here's one way to expand the idea, adapted from Larson and Segal 1995, pp. 74-6. Following Evans 1976, let's call a conditional of the form 'If A is true, then B is true' a validating conditional. And let's say that an inference from A to B is truth-theoretically valid according to a particular truth-theory iff one can derive its validating conditional using the truth-theory and logic.⁴³ So, for instance, the inference from 'Emma is walking slowly' to 'Emma is walking' is truth-theoretically valid according to an eventish truth-theory, because, from the truth-theorems for each, one can use logic to derive the validating conditional that, if the former is true, then the latter is true. Thus, one way a truth-theory can be said to capture a validity is for the inference to be truth-theoretically valid according to the truth-theory. Now, the worry might be formulated as follows: a truth-theory's capturing a validity is deprived of explanatory power if the truth-theorems deployed in deriving the validating conditionals are not (or need not be) true.

To assess the worry, we need a better sense of just what is supposed to explain what. One possible claim, for instance, is that an account of semantic competence ought to, or can, explain why speakers find such inferences obviously valid. Here the explanans is that we cognize a certain truth-theory, one with certain truth-theoretic validities (and that we are capable of reasoning logically with the truth-theory's results); and the explanandum is our judgment concerning those inferences which are truth-theoretically valid according to the cognized truth-theory. It's unclear, however, why this goal is threatened by the possibility of untrue truth-theorems—or at least truth-theorems untrue

⁴³ We need to explicitly allow the resources of logic, because the truth-theory's production rules need not—and for Larson and Segal and for Ludlow do not (see below)—include the full resources of logic. (I should make clear that Evans' concern—to

owing to excess ontological commitment. Why can't one explain speakers' finding obvious the inference from 'Emma is walking slowly' to 'Emma is walking' by adverting to what is logically implied under certain suppositions we, or our semantic modules, make? There's no trouble seeing why the logical relations cognized truth-theorems' right-hand-sides bear to one another might be relevant to speakers' finding these inferences obviously valid. Evidence pertaining to inference can thus remain relevant in discriminating among competing, not necessarily true truth-theories.⁴⁴

8) Perhaps, though, it will be maintained that the truth-theory's capturing an inference ought to or can explain, not (or, not just) why speakers find the inference compelling, but why the inference is valid. This might seem to ask too much of the truth-theory. For among the truth-theoretic validities, on any plausible truth-theory, are logical validities; and it can be questioned whether one has really supplied an explanation of the validity of, say, conjunction elimination, if conjunction elimination is used in the derivation of its validating conditional.⁴⁵ Still, one might hold that being truth-theoretically valid according to the appropriate truth-theory provides an explanation of the validity of certain inferences—viz., those not deployed in deriving their validating conditionals. What would do the explaining in this case would not be the speaker's cognizing the truth-theory, but rather the fact that these inferences are truth-theoretically valid according to this truth-theory—otherwise, psychologism looms. However, given

reject a similar Davidsonian attempt to explicate *structural* (as opposed to lexical) validity—is not mine here.)

⁴⁴ One might worry that there is a problem of explanatory over-generation here. Why don't I find obvious the inference from 'Emma is walking slowly' to 'There is a walking'? But if there is a problem, it arises equally for ontologizing semanticists who would thus explain these aspects of our inferential competence.

this explanation of the relevant validities, one might then explain speakers' grasp of them by appealing to their cognizing the truth-theory that explains the validities themselves.

Here we find ourselves in a position parallel to the one we were in above when considering the meaning-stating project. We must ask whether a truth-theory invoked as part of an account of semantic competence—one, moreover, that can apparently help account for aspects of inferential competence as well—must also be what explains validity (or at least this range of entailment facts). Different inferences will be truth-theoretically valid according to different truth-theories. So, not just any truth-theory will explain the validity of the relevant inferences that are in fact valid. Why must it be that the truth-theory that does, supposing one does, is distinguished in part by being the one cognition of which helps explain semantic competence?

As before, an ontologizing semanticist pressing the current line of defense commits herself to several substantial claims. First, she commits herself to this explanatory project's legitimacy against those who would question whether there's a semantic explanation of this class of validities at all (cf. Brandom 1994). Second, she commits herself to the explanation taking the truth-theoretic form just outlined. And, third, she commits herself to the convergence of the truth-theory that explains such inferences with that cognized by speakers. That is, even if there *is* a truth-theoretic explanation of this particular class of validities, it's a further claim to hold that the truth-theory that provides this explanation is the one whose cognition accounts for speakers' semantic competence (and perhaps also aspects of their inferential competence).

⁴⁵ Such explanations needn't clash, however, with more standard characterizations of logical validity such as that in terms of preservation of truth in all models, since the latter's concern is which validities are logical validities.

9) One might argue that it's the truth-theory's role in a speaker's larger cognitive economy that requires that its theorems be true. Such an argument might take various forms, but let me sketch just one. The argument uses two premises: first, that the Language of Thought hypothesis is true; and, second, that the role of a truth-theory is to translate between natural language sentences and representations in the Language of Thought. According to the first premise (irrelevant qualifications aside), it's nomologically the case that one believes that P if and only if one stands in some relation R to a mental representation (a sentence-like token realized in the brain) with the content that P—similarly, for other propositional attitudes. In particular, one understands someone as having said that P if and only if one stands in some relation R' to a mental representation with the content that P, and one believes what the person said if and only if one stands in R to that mental representation. It follows from the second premise that it's nomologically the case that, bracketing any pragmatic effects, one accepts what someone says in uttering S if and only if one stands in R to a mental representation characterized by the right-hand-side of the truth-theorem for S that is yielded by the truth-theory one cognizes.

The rough picture—assuming sequential processing just to keep the picture simple—is this. When you hear an utterance to which you succeed in assigning some syntactic representation(s), the relevant syntactic representation gets fed to your semantic module, which then assigns it truth-conditions. The module's job is then to output a mental representation of these truth-conditions. Bracketing further pragmatic processing, this mental representation is then treated as what the producer of the utterance said. If you are then inclined to accept on testimony what was said, then this mental representation is

placed “in the belief-box”—that is, it is placed in the relevant relation R to you. Similarly, if there’s a mental representation in your belief box that you wish to communicate, it gets sent to the semantic module where it’s associated with a syntactic form, the idea being that the mental representation has a format found on the right-hand-side of some truth-theorem. The syntactic form is then associated with a phonological form to yield something utterable.

Suppose, though, that the cognized truth-theory may be false. Then it’s possible to believe something that’s true by standing in R to a mental representation with false content. But then the mental representation can’t have the *same* content as what is believed. So, the Language of Thought Hypothesis and the more specific claim concerning the role of the cognized truth-theory can only be accepted if it’s also held that the cognized truth-theory is true.

There are several possible responses to this line of thought. Most obviously, its two assumptions are controversial. First, there is the Language of Thought hypothesis itself. Cognitivist neo-Davidsonianism is an hypothesis about what attitudes human speakers (tacitly) have; the Language of Thought hypothesis concerns what’s involved in having an attitude. To endorse the former is not yet to endorse the latter. It’s of course a large question whether the Language of Thought hypothesis should be endorsed. Second, the idea that the semantic module has the job of translating between one’s idiolect and the Language of Thought in the fashion sketched above is quite speculative. Even if the cognized truth-theory does yield a Mentalese mental representation in response to the uttered sentence, it doesn’t follow that *that* representation is what goes in the belief box. It’s possible, for example, that that mental representation must undergo further (non-

pragmatic) processing before being placed in the belief box, perhaps including removal of its particular existential commitments—so that what goes in the belief box is as it were further downstream, from a processing perspective, than what appears on the t-theorem’s right-hand-side. It’s also possible that what goes in the belief box is something *upstream* from the t-theorem’s right-hand-side: for example, perhaps it’s the left-hand-side’s syntactic representation itself that’s the Mentalese representation destined for the belief box, rather than the interpretation of it supplied by the semantics module.⁴⁶

Finally, it may be possible for someone resisting the ontologizing semanticist to accept the gist of the argument by modifying its details. Specifically, much of the motivation for the Language of Thought hypothesis is preserved if we substitute for the particular formulation provided above the claim that it is nomologically necessary that you believe that P iff you stand in some relation R to a mental representation with a content *appropriately related* to the claim that P. This weaker requirement might be satisfied by an existentially committing claim, believed under a useful existential supposition. The intuitive idea is that our sub-personal cognitive system as a whole, or at least more than just our semantic module, might make as if certain things exist.⁴⁷

To conclude this section, let me sum up what the ontologizing semanticist would need to show if she were to hope to deploy successfully some positive version of the

⁴⁶ This would accord with Ludlow’s (1999, pp. 22-6 and 165-9) hypothesis that I-language is the Language of Thought. Cf. Carruthers 2002.

⁴⁷ What might it be for a mental representation to be appropriately related to a claim with content P? Perhaps that it be paired by the truth-theory cognized by the semantic module with a sentence that in fact means that P. So, when the relevant mental representation quantifying over events is in your belief box, you believe that Emma is walking slowly—because that mental representation is paired with [Emma is walking slowly] by the truth-theory.

belief strategy. On either the speaker or theorist version, it must be the case that apparently existentially quantifying right-hand-sides really are ontologically committing, and not instead canonically rendered in some ontologically deflationary way. In order for the speaker version to work, the ontologizing semanticist must in addition show that speakers do in fact believe—in an ontologically committing way—the truth-theories ascribed to them and that they believe, not only the relevant mundane belief, but also that its appropriate syntactic representation is true. In order for the theorist version to work, she needs to show that theorists who themselves believe the truth-theories they ascribe are justified in doing so on the basis of their ascription. That is, it must be that an accurate ascription of a truth-theory as part of an account of semantic competence requires the truth-theory's truth. Showing this might be attempted in a variety of ways, each with their own substantial commitments.

V. The Canonical Notation Strategy

I turn now to the second broad strategy the ontologizing semanticist might deploy: the canonical notation strategy. Recall that this strategy starts from the thought that sentential surface structure is not necessarily a good guide to ontological commitment. To determine someone's ontological commitments, it is necessary that her beliefs—or, the sentences expressive of her beliefs—first be represented in the proper way. The strategy's central claim, then, is that how the beliefs ought to be represented for ontological purposes is a function of the truth-theory ascribed in an account of the subject's semantic competence. On this approach, it suffices for ontological commitment that the speaker's antecedently held beliefs when thus canonically represented logically imply existential

claims—regardless of whether the speaker or anyone else *believes* the ascribed truth-theory.

We need now to see how the ascription of a truth-theory might play a role in assigning a canonical representation. Two main ways can be distinguished, according to which side of a sentence's truth-theorem is drawn upon. About both we can be somewhat brief: the left-hand-side approach because it simply does not jibe with the actual practice of ontologizing semanticists; the right-hand-side approach because some of the relevant issues recur from the discussion of the theorist's belief strategy.

A. Left-Hand-Side

A brief reminder of one strand in philosophico-semantical history—from Quine, through Davidson, to our neo-Davidsonians—provides an approach to the first possibility.

For Quine, the proper way to represent a claim, if one's concern is ontological commitment, is by associating the claim with a sentence in a formal language whose form makes explicit this formula's logical relations to other sentences of the formal language. On Quine's view (e.g., 1953), the formal language deserving the title 'canonical notation' is that of classical first-order logic. There's no suggestion that the canonical representation associated with a claim has any psychological reality in the sense of playing a role in an empirical theory of semantic (or, for that matter, logical) competence. A logical form can be thought of as a replacement of a natural language sentence for certain purposes. Because it wears its logical relations to other formal sentences on its sleeve, it can be used in particular to determine ontological commitment,

since such commitment consists in its logically implying an existentially quantified sentence.

Davidson follows Quine in awarding classical first-order logic the title of ontologically revealing canonical notation, albeit with more of an eye than Quine on capturing what logical form natural language sentences in fact have (as opposed to what logical forms a theorist can for certain purposes associate them with). There is an interpretive question as to whether Davidson himself thinks of a sentence's logical form as what its t-theorem contains on its right-hand-side (thus as something an application of the truth-theory to the sentence reveals about it) or whether his analyses of logical form provide rather a sentence's proper object-level representation for truth-theoretic purposes—that is, provide that to which the truth-theory is applied to yield truth-conditions.⁴⁸ But it's the latter construal that's relevant to the historical strand I'm here following. For example, his eventish analysis of adverbial modification can and has been taken to supply the underlying logical forms for adverbially modified English sentences, for which one should then state truth-conditions. Those truth-conditions will themselves quantify over events; but what's notable on this line is that event-places are as it were already supplied in the left-hand-side descriptions of the target sentences (cf., e.g., Lycan 1984).

The Davidsonian idea is brought closer to current neo-Davidsonianism by the suggestion—found, for example, in Harman 1972—that linguists empirically investigating the syntactic features of natural languages might benefit from positing

⁴⁸ For example, Davidson 1977, pp. 209-10, suggests that, although it's a “convenience” if we can “find” a canonical notation in the object language and use it to bring logical form to the “surface” so that one may supply a “transparent” description of the sentence's structure on the left-hand-side, this is not necessary.

logical form, so construed, as among the structures needed to explain their data. So, the hypothesis is that logical form is (or is among) the syntactic representation(s) language-users assign natural language sentences; and, what's more, that it is to *this* representation that speakers then assign semantic properties via a cognized truth-theory for the language.

There is thus an historical lineage connecting Quinean canonical notation (with its claimed ontological significance) to the syntactic representations found on the left-hand-side of the t-theorems ascribed in cognitivist accounts of semantic competence. Our lightning tour thus provides some reason for thinking that the canonical notation strategy might be what at least some ontologizing semanticists have in mind, with these left-hand-side syntactic representations playing the role of canonical representations.⁴⁹

The relevant conception of structure, however, has evolved along with the conception of its role in theorizing. Though the semantically relevant syntactic representations get called 'LF'—short for logical form—they bear little resemblance to the formulas of 1st-order logic. There are a variety of reasons for this—for example, because, as syntactical representations of natural language sentences, LF's are called upon to help explain judgments concerning the (un)grammaticality of various distributions of expressions, whereas the formulas of 1st-order logic are constructed with an eye toward making explicit certain deductive relations among them. It doesn't follow, however, that there couldn't be one structural representation fit for both of these roles. Perhaps, at the end of the day, this will be LF—even if it contains some structural

⁴⁹ It might be objected that this would award syntax a role in settling ontological disputes, not accounts of semantic competence. But since the relevant syntactic representation is the one fed into the semantic theory—that is, is the interface between syntax and semantics—evidence from the latter has a crucial bearing on the nature of the former.

features irrelevant from a logical point of view and some irrelevant from a grammatical point of view. It's not the common practice of syntacticians, semanticists, or logicians to provide inference rules for LF's. But there's no reason one couldn't try. And indeed some have. In fact, a promulgator of this project—the project of “natural logic”—is Ludlow 2002. If the project could be carried through, LF would indeed be a logical form, even if it differed from the logical forms found in 1st-order logic.

But it must be emphasized that there's no a priori reason why the kinds of structural features important to the study of syntax, semantics, and logic, respectively, should either coincide or be combined into one representation in order to account for cognition (cf. Larson and Segal 1995, pp. 74-6). The data these inquiries attempt to explain, in part by positing structural features, differ markedly. Syntax needs to explain, again, judgments concerning the (un)grammaticality of various distributions of expressions; semantics our comprehension of expressions; logic the inferential relations among them. Why think that our cognitive economy in fact utilizes one structural representation to encode the various features relevant to these disparate tasks (even if they in various ways constrain one another)? Of course, we are granting the empirical credence of cognitivist truth-theoretic accounts of semantic competence, and so in particular the claim that there is a level of representation that provides the interface between syntax and semantics—that is, that encodes the syntactic features relevant for the assignment of semantic value by a truth-theory. But this leaves open the empirical question whether there's compelling reason to identify such structural representations with some conception of logical form—and this is less clear. That it is an empirical question—one for which a positive answer must currently count as “highly speculative”

(Ludlow 1997, p. 204)—is what I want to emphasize. Thus the label ‘LF’ is often used, not so much as an abbreviation, as out of caution.

In any event, it is clear that not all ontologizing semanticists tie their claims to the empirical fortunes of such speculation. The left-hand-side canonical notation strategy only gets off the ground in a given case, of course, if the relevant structural representation does logically imply the existence of *k*'s. But Ludlow and Larson and Segal introduce events into the semantics of adverbs without introducing events into their syntax (see Ludlow 1999, p. 73, and Larson and Segal 1995, p. 472). There is thus no (syntactic) quantification over events in a sentence like [Emma is walking slowly]; nor is any suggestion made that this LF logically implies an LF that does contain such quantification—whatever “natural logic” may someday yield. But Ludlow and Larson and Segal see such semantic theorizing as ontologically committing, given the mundane belief that Emma is walking slowly. They must therefore envision a different argument than the left-hand-side canonical notation strategy.⁵⁰

I thus pass to the second way of pursuing the canonical notation strategy. It should be noted, however, that the questions we will raise about it could be raised as well for versions of the strategy we have just discussed. So, even if it were the practice of ontologizing semanticists only to draw ontological conclusions based on the semantically

⁵⁰ That there's no syntactic quantification over events in [Emma is walking slowly] doesn't preclude that there is in other syntactic forms. But Larson and Segal and Ludlow take the treatment of such sentences alone as sufficient for the ontological conclusions. Noting their presence in other syntactic forms would thus leave unanswered why they hold that adverbial modification yields ontological commitment to events. (Of course, they might be wrong about the need for event places in the syntactic analysis of adverbial modification. See Higginbotham 2000, p. 70, for example, for arguments that these event positions must be “visible” to syntax.)

relevant syntactical representation of a sentence, there would still be further questions for such a strategy to address.

B. Right-Hand-Side

Suppose, then, it is claimed that it's the right-hand-side of a truth-theorem that provides the canonical representation of a sentence for the purpose of determining ontological commitment. The fundamental question that must be asked is: why think that? And, in particular, why think that the role of a truth-theory in accounting for semantic competence establishes or supports this?

The three most natural reasons, I think, whether offered separately or together, are these: first, because the right-hand-side provides the target sentence's logical form; second, because it states what the sentence means; and/or, third, because it expresses what is believed by a person who assents to the target sentence construed literally. About each, there are two questions that must be asked: (1) Does the right-hand-side in fact deliver these goods? And (2) If it does, does its doing so in fact bestow upon it this privileged ontological role?

It's perhaps easiest to answer question (2) in the affirmative if the right-hand-side provides logical form. The logical form of sentence is a representation that makes formally explicit its logical relations to other logical forms. In particular, it makes formally explicit whether it logically implies any existentially quantified claims. Given the Quinean conception of ontological commitment, it's thus clear why someone might

hold that a sentence's logical form should constitute its ontologically canonical representation.⁵¹

It's not as obvious how one should address question (2) having established, or at least claimed, that the right-hand-side rather states what the target sentence means and/or expresses what someone who assents to it believes. Suppose the target sentence is 'It's contingent that Bill is tall' and that the cognized truth-theory assigns as truth-conditions: there's a possible world in which Bill is not tall. Suppose further that this is what the target sentence means. One might hold that it's also the case that the target sentence means that it's contingent that Bill is tall—not because the target sentence is ambiguous, but because these two statements of meaning agree in content. But which is privileged for the purposes of determining ontological commitment?⁵² The ontologizing semanticist

⁵¹ The semantics module, according both to Larson and Segal and to Ludlow, employs what may seem a rather restricted range of production schemata. Larson and Segal include Universal Instantiation, Substitution of (Provable) Equivalents, and Substitution of Identicals; Ludlow includes the latter two plus a schema for reducing redundancies. (The differences reflect the different fragments of English they consider.) It's not just that the full power of even propositional logic is not needed: it's by adverting to the hypothesis that the semantics module employs relatively weak schemata that they avoid problems of over-generation (Larson and Segal 1995, pp. 34-7, and Ludlow 1999, p. 217, n. 9). (Larson and Segal 1995, pp. 568-70, n. 2, also advert to this hypothesis to help handle empty names.) The production schemata, they say, are indeed not to be thought of as inference rules at all (Larson and Segal 1995, p. 559, n. 13, and Ludlow 1999, p. 35). It might therefore be objected that forms assigned on the right-hand-side aren't logical forms. This doesn't follow, however. For it could be that, although the semantics module doesn't exploit the full resources represented in the right-hand-sides, they are nonetheless there, perhaps even to be exploited elsewhere in our cognition (by a logic module, if you will—cf. Larson and Segal 1995, p. 75, but also for criticism Szabó 1997).

⁵² This is simply a version of the objection Alston 1958 raises against Quine: if A and B are synonymous or otherwise suitably related, why view A as ontologically revelatory of B rather than the other way round? The objection, however, doesn't have force against Quine. He held that questions of ontology simply don't arise in ordinary language (Quine 1981, p. 9: "a fenced ontology is just not implicit in ordinary language"); and of course he countenanced no relation of synonymy—in particular, he didn't maintain that the pragmatically guided (that is, holistic, empirically guided) replacement of unregimented sentences by canonical ersätze must preserve "meaning". Cf. Quine 1960, Chapter II.

pursuing this line, unless she denies the claim of synonymy in the first place, must hold that the meaning-statement drawing upon the ascribed truth-theory's right-hand-side—so, in the case at issue, the one that quantifies over possible worlds—provides the canonical representation for purposes of determining ontological commitment. But what might her grounds be? This is the question with which we started. So, the claim that a truth-theorem's right-hand-side states what the target sentence means, even if true, does not provide an answer to this question if there may also be other statements of what the sentence means. Parallel remarks apply to the claim that the right-hand-side expresses what someone who assents to the target sentence believes.⁵³

There is another issue, however, that can be raised with respect to all three versions of the right-hand-side strategy. As with the various versions of the belief strategy, we can ask whether the quantifiers on the right-hand-side are correctly construed as ontologically committing—for example, whether they are objectual or substitutional. The Quinean idea is to render the sentences in such a way that this question simply doesn't arise. But that means, for Quine, either that our canonical notation shouldn't contain such alternative quantifiers or that, if they do, they must be clearly distinguished so as to avoid ambiguity. If alternative quantifiers can occur in canonical notation, but are notationally distinguished, then the question of how to

⁵³ Note that the Alston question can be reasonably rebuffed by someone claiming that the right-hand-side provides the target sentence's logical form. Suppose, for instance, that the logical form of some target sentence involves existential quantification, but that it is also synonymous with—or in some other way suitably related to—some sentence that does not. The Alston question is: why should we take the existentially quantified sentence to reveal the other sentence's ontological commitment instead of taking the other sentence to show how to eliminate the existentially quantified sentence's apparent ontological commitment? The answer is simply that the existentially quantified sentence, because it provides the target sentence's logical form, makes explicit what the sentence implies in a way the other sentence doesn't.

construe some quantifier occurring in canonical notation might not arise. But the question can indeed arise of how best to represent some natural language sentence in canonical notation in the first place so far as quantificational structure is concerned: should its ersatz contain objectual or substitutional quantifiers (or neither or both)? This is true as well for the ontologizing semanticist. There are of course empirical demands on her project that don't bear on Quine's (and, in any event, weren't accepted by him as legitimate empirical demands on any project); so, she must be responsive to different kinds of reasons in addressing whether substitutional quantifiers appear on the right-hand-sides of t-theorems. But it is a question she must address.⁵⁴

What of question (1): does the right-hand-side in fact provide the target sentence's logical form, state its meaning, or express what is believed by someone who assents to it? Here we encounter questions quite similar to some already addressed in discussing the theorist's belief strategy.

There's certainly room to object that the right-hand-side provides logical forms. Suppose you deny that there are events but do think that Emma is walking slowly. Then you have reason to deny that it logically follows from the fact that Emma is walking slowly that there are events. You therefore have reason to deny that the right-hand-side standardly assigned by a truth-theory for the sentence [Emma is walking slowly] provides its logical form.

It will be recalled that earlier we placed to one side moves that would avoid ontological commitment by revising a truth-theory in light of antecedently held metaphysical views. So, it's important to note that this is not what is going on here. The

⁵⁴ There remains, as remarked above, the question of whether substitutional quantifiers are devoid of ontological commitment.

objector does not question ascription of the truth-theory; rather, she questions, on the basis of her ontological scruples, only that the right-hand-side provides the target sentence's logical form. She is suggesting that for a truth-theory to be correctly ascribed—that is, for cognition of it to explain in part our semantic competence—it's not necessary that it assign logical form.

To rebuff this objection, an argument is needed for why the right-hand-side must be construed as providing the target sentence's logical form. We have already encountered two reasons one might offer. First, it might be suggested that the ascription of the truth-theory explains also aspects of our inferential competence. We saw, however, that this could be the case even if the truth-theory ascribed isn't true. But if the truth-theory needn't be true, then the right-hand-sides can't be considered logical forms, at least in any sense that would privilege them in determining ontological commitment, since their logical relations won't exactly capture those of the target sentences. Second, it might be suggested that the truth-theory explains, not aspects of our inferential competence, but certain validities themselves. Adverting to such a claim in the context of advancing this version of the canonical notation strategy requires taking on the same commitments discussed above in the context of the theorist's belief strategy: a commitment to this explanatory strategy and to its convergence with the project of accounting for semantic competence.

Parallel remarks apply to the other two suggestions as well. As we've seen, it's not obvious that a truth-theory ascribed to account for semantic competence must also state or otherwise show what sentences mean. To hold otherwise commits the ontologizing semanticist to substantial claims concerning the meaning-stating project and

its relation to the project of accounting for semantic competence. Nor is it obvious that the right-hand-sides express what speakers assenting to the target sentences believe. For example, as discussed, it's a substantial claim that the right-hand-sides indicate how a speaker assenting to the target sentences represents her belief in the Language of Thought.⁵⁵

Thus, in addition to the issues raised in discussing question (2), the most natural ways of running the right-hand-side canonical notation strategy require substantial commitments of the ontologizing semanticist, commitments like those involved in running certain versions of the theorist's belief strategy. Ontologizing semanticists may well find (some of) those commitments in any event attractive. But it's unclear that endorsing a cognitivist truth-theoretic account of semantic competence commits you to them.

VI. Concluding Remark

I have surveyed some of the further commitments an ontologizing semanticist might have to make in order to secure her form of argument—at least in drawing positive ontological conclusions. Perhaps, in the eyes of some, the survey succeeds, not in problematizing the move from an account of semantic competence to ontology, but rather in helpfully pointing out the direction(s) one should look for the further details that secure it. But it should be recalled that we began by noting that it's prima facie surprising that such contingent features of human cognition might bear on matters of such metaphysical

⁵⁵ I have not discussed an alternative way one might develop this particular line: viz., by holding that the right-hand-sides indicate the structure of the proposition (understood as a structured, but non-linguistic, abstract object) believed by someone who assents to the target sentence. This line would clearly involve further commitments of its own.

moment. It's hard enough to pull off transcendental arguments adverting to features necessary to any mind or representational system: their empirical cousins face even further obstacles and must bear a heavy burden of proof.⁵⁶

References

- Alston, William. (1958) "Ontological Commitment," Philosophical Studies 9.
- Brandom, Robert. (1994) Making It Explicit (Harvard).
- Carruthers, Peter. (2002) "The Cognitive Functions of Language," Behavioral and Brain Sciences 25.
- Carston, Robyn. (2002) *Thoughts and Utterances* (Blackwell).
- Chomsky, Noam. (1975) Reflections on Language (Fontana/Collins).
- Davidson, Donald. (1967a) "Truth and Meaning," Synthese 17.
- Davidson, Donald. (1967b) "The Logical Form of Action Sentences," The Logic of Decision and Action, ed. N. Rescher (Pittsburgh).
- Davidson, Donald. (1973) "Radical Interpretation," Dialectica 23.
- Davidson, Donald. (1977) "The Method of Truth in Metaphysics," Midwest Studies in Philosophy 2. Cited as reprinted in his Inquiries into Truth and Interpretation (Oxford, 1984).
- Davidson, Donald. (1983) "A Coherence Theory of Truth and Knowledge," Truth and Interpretation, ed. E. Lepore (Blackwell).
- Evans, Gareth. (1976) "Semantic Structure and Logical Form," Truth and Meaning, eds. G. Evans and J. McDowell (Oxford).

⁵⁶ Material related to this paper was presented at the University of London, the University of Pennsylvania, and the University of Maryland. I thank audience members and this

- Field, Hartry. (1980) Science Without Numbers (Princeton).
- Gross, Steven. (2002) "Putnam, Kontext, und Ontologie," *Hilary Putnam und die Tradition des Pragmatismus*, eds. M. Raters and M. Willaschek (Suhrkamp).
- Gross, Steven. (forthcoming) "Context-Sensitive Truth-Theoretic Accounts of Semantic Competence," *Mind and Language*.
- Gross, Steven. (submitted) "The Biconditional Doctrine: Contra Kölbel on a 'Dogma' of Davidsonian Semantics."
- Harman, Gilbert. (1972) "Deep Structure as Logical Form," Semantics of Natural Language, eds. D. Davidson and G. Harman (Reidel).
- Higginbotham, James. (1985) "On Semantics," Linguistic Inquiry **16**.
- Higginbotham, James. (2000) "On Events in Linguistic Semantics," Speaking of Events, eds. J. Higginbotham, F. Pianesi, and A. Varzi (Oxford).
- Hume, David. (1779) Dialogues Concerning Natural Religion.
- Kölbel, Max. (2001) "Two Dogmas of Davidsonian Semantics." Journal of Philosophy **98**.
- Larson, Richard, and Gabriel Segal. (1995) Knowledge of Meaning (MIT).
- Ludlow, Peter. (1997) "Logical Form and Grammatical Form: Introduction," Readings in the Philosophy of Language, ed. P. Ludlow (MIT).
- Ludlow, Peter. (1997a) "Preface," Readings in the Philosophy of Language, ed. P. Ludlow (MIT).
- Ludlow, Peter. (1999) Semantics, Tense, and Time: An Essay in the Metaphysics of Natural Language (MIT).

journal's referees for helpful comments. I also thank Peter Ludlow for encouragement.

- Ludlow, Peter. (2002) "LF and Natural Logic," Logical Form and Language, eds. G. Preyer and G. Peter (Oxford).
- Lycan, William. (1984) Logical Form in Natural Language (MIT).
- McTaggart, J. M. E. (1908) "The Unreality of Time," Mind **68**.
- Parsons, Charles. (1971) "A Plea for Substitutional Quantification," Journal of Philosophy **68**.
- Parsons, Terence. (1990) Events in the Semantics of English (MIT).
- Pietroski, Paul. (2003) "The Character of Natural Language Semantics," *Epistemology of Language*, ed. A. Barber (Oxford).
- Quine, W. V. O. (1948) "On What There Is," Review of Metaphysics **2**.
- Quine, W. V. O. (1953) "Mr. Strawson on Logical Theory," Mind **62**.
- Quine, W. V. O. (1960) Word and Object (MIT).
- Quine, W. V. O. (1970) "Methodological Reflections on Current Linguistic Theory," Synthese **21**.
- Quine, W. V. O. (1977) Review of Truth and Meaning: Essays in Semantics, eds. G. Evans and J. McDowell (Oxford, 1976), in Journal of Philosophy **74**.
- Quine, W. V. O. (1981) "Things and their Place in Theories," in his Theories and Things (Harvard).
- Soames, Scott. (1989) "Semantics and Semantic Competence," Philosophical Perspectives **3**.
- Stalnaker, Robert. (1974) "Pragmatic Presuppositions," Semantics and Philosophy, eds. M. Munitz and P. Unger (NYU).

Szabó, Zoltán Gendler. (1997) Review of Larson and Segal (1995), The Philosophical
Review **106**.