

What Meanings Probably Are Not (Given Autonomous Syntax)

[draft 31Aug04; references incomplete]

Paul M. Pietroski, Dept. of Philosophy & Dept. of Linguistics, Univ. of Maryland

Davidson (196x, 196x, 196x) offered a famous conjecture that proved enormously fruitful: for each natural language L, there is a theory of truth for L than can serve as an adequate theory of meaning for L. Given developments of Davidson's seminal proposals, and those of Montague (1970, 1974), it now seems clear that the Frege-Tarski toolkit—originally designed for the invention and study of formal languages—can be usefully applied to the study of natural languages. Only intellectual curmudgeons refuse to see the successes of “truth-conditional” semantics. But the program also faces deep difficulties, not merely an absence of plausible accounts for various constructions. Critics often focus on a cluster of considerations associated with (late) Wittgenstein, Austin, Strawson, and Chomsky: in light of all the things that speakers can do with language, and the ubiquitous sensitivity of truth to communicative situation, it seems implausible that words have the “satisfaction conditions” required by empirically adequate theories of truth (that can also serve as explanatorily adequate theories of meaning).¹ In this paper, I focus on a different though potentially supporting line of thought, arising from a theoretically attractive conception of human language shared by some friends of truth-conditional semantics.

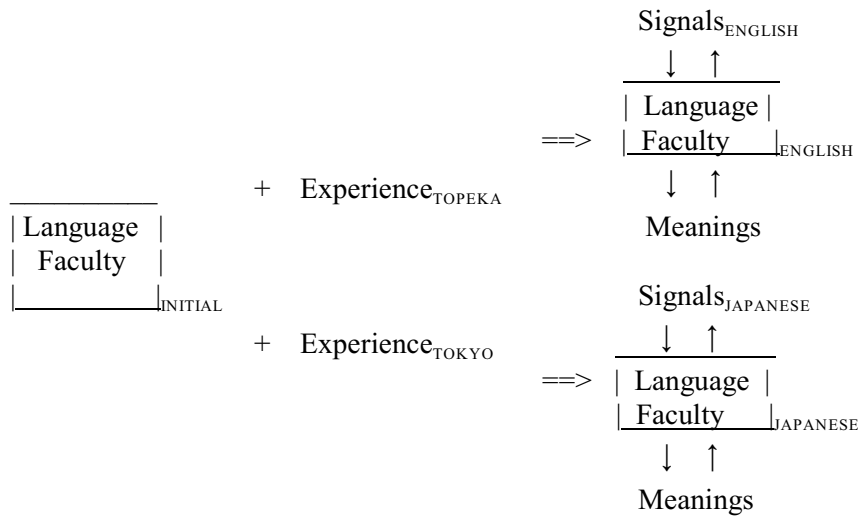
If we follow Chomsky in taking the study of natural language to be at least largely a study of the human “language faculty,” then discoveries associated with the “autonomy of syntax” thesis (and poverty of stimulus arguments) should make us skeptical of many claims about natural language semantics—including the claim that there are theories of *truth* for natural languages, as opposed to theories of certain intrinsic features of linguistic expressions that *constrain without determining* how those expressions can be used to make assertions. In short, given that human biology imposes substantive constraints on how natural language can be understood, many current conceptions of meaning are implausible; in particular, meanings are unlikely to be functions from contexts to Truth/Reference/Satisfaction (TRS) conditions.

1. Some Terminology

Let's say that a *language* is any system for associating signals with meanings, whatever those are. A *natural language* is one that normal human children can acquire in a developmental setting that is not atypical for members of our species. In so far as ‘system’ is ambiguous—say, as between a set of abstract rules, and a (presumably physical) device that instantiates such rules (as an idealization)—‘language’ is also ambiguous. So following Chomsky (1986), let's use the technical term ‘I-language’ to talk about the relevant aspects of human minds/brains. Each speaker of a natural language has an I-language, by virtue of which she can associate endlessly many linguistic signals with meanings in a certain way. This leaves open all the interesting questions. We can ask to what degree, and in what respects, I-languages are: determined by human biology, as opposed to experience; “encapsulated” from other aspects of human

psychology; shared by speakers within a given linguistic community; etc. But at least we can address such questions free of the confusion that postulating grammatical principles cannot really be a part of investigating human psychology.

Idealizing, we can think of children as coming equipped with a language faculty, whose initial state can be changed (through experience and growth) within constraints that we can try to discover; where this faculty typically settles into one or more stable adult states, each of which is an I-language.² The picture is familiar.



Unsurprisingly, the utility of this idealization depends in part on the answers to the questions it lets us ask. If each I-language of each speaker turns out to be an idiosyncratic interaction effect—involving most of that speaker’s central nervous system, and depending heavily on her particular course of experience—then we should probably look elsewhere for interesting and theoretically tractable linguistic phenomena. But without trying to defend specific answers in detail, I want to rehearse a few well-known arguments for the Chomskyan conception of the language faculty that many of us endorse. For in my view, these considerations fit ill with the hypothesis that I-languages associate linguistic signals with (functions from contexts to) TRS conditions. If investigation reveals that I-languages are subject to severe constraints imposed by the nature of the language faculty, then we should ask whether these constraints are plausibly characterized as limitations on how I-languages can associate signals with TRS conditions; and likewise for other conceptions of meaning according to which I-languages associate signals with aspects of the mind/language-independent environment. My claim, following Chomsky (200x), will be that such conceptions of the language faculty are implausible.

2. Some Facts

In this section, I illustrate some constraints respected by I-languages, focussing on cases that should be familiar and relatively uncontroversial. More tendentiously, I will assume (without argument here) that these elementary cases are indicative of a more general point about the last fifty years of linguistics: the theoretically interesting phenomena are typically “negative,” concerning possible structures not exhibited in natural language and possible meanings not expressible with actual expressions. This claim is, of course, of a piece with Chomsky’s idea that questions about language acquisition—and poverty of stimulus arguments that suggest nativist answers—should occupy center stage in theoretical investigations of human language. In sections three, I’ll argue that by (wisely) adopting a conception of the language faculty as the source of the negative facts in question, one thereby renders suspect various “externalist” conceptions of meaning. But first, the review.

Consider the following sequence of lexical items: *hiker, lost, kept, walking, circles*. This string of words might well prompt the thought indicated with (1), as opposed to the thought indicated with (2).

- (1) The hiker who was lost kept walking in circles
- (2) The hiker who lost was kept walking in circles

While (2) is a perfectly good English sentence, which can be used to express a coherent thought, it is unexpected. So it is striking that (3) unambiguously indicates the yes/no question corresponding to (2).

- (3) Was the hiker who lost kept walking in circles

We hear (3) as synonymous with (4), not (5).

- (4) Was it the case that the hiker who lost was kept walking in circles
- (5) Was it the case that the hiker who was lost kept walking in circles

Likewise, (8) is the yes/no question corresponding to the bizarre (7), not (6).

- (6) The child who was fed the soup fed the kitten.
- (7) The child who fed the soup was fed the kitten.
- (8) Was the child who fed the soup fed the kitten?

And we know that natural language is not hostile to ambiguity, given examples like (9).

- (9) The lawyers who can duck and hide whenever visiting relatives may scare them
saw every doctor who lost her patience with patients.

If a string of words can be understood as a sentence, but it cannot be understood as having a sentential meaning easily expressed with a good sentence formed from those words, then this negative fact calls for explanation—especially if the actual meaning is more “cognitively surprising” than the nonmeaning. (See Chomsky [1965], Higginbotham [1985].) Standard explanations for the nonambiguity

of (3) posit a constraint that (one way or another) precludes extraction of auxiliary verbs from relative clauses; see Ross (196x), Travis (198x). One hypothesizes that the transformation indicated in (3b) is licit, while the transformation indicated in (3a) is not.

- (3a) Was {[the [hiker [who ___ lost]_{RC}]][kept walking in circles]}
 |_____*_____|
- (3b) Was {[the [hiker [who lost]_{RC}]][___ kept walking in circles]}
 |_____|

If a string of words fails to have *any* coherent interpretation, that is a special case of nonambiguity. But there are many ways in which such failure can come about, some more interesting than others.

For example, while (10) and (11) are fine, (12) and (13) are defective; though (13) is closer to word-salad than (12).

- (10) The hiker who was lost whistled. (12) *Was the hiker who lost whistled?
 (11) A vet saw the dog that was found. (13) *Was a vet saw the dog that found?

We can start to explain this by noting that (12) would be the word-string corresponding to two transformations, shown in (12a) and (12b);

- (12a) Was {[the [hiker [who ___ lost]_{RC}]][whistled]}
 |_____*_____|
- (12b) Was {[the [hiker [who lost]_{RC}]]*[___ whistled]}
 |_____|

where the asterisk in (12b) indicates the anomaly of asking whether or not the *hiker* was whistled—cf. ‘Was the song whistled’. The defectiveness of (12b) is not so severe that it keeps us from hearing (12) as a question, albeit an odd one, that we could echo. (The hiker who lost was whistled?) But the structural violation in (13a) evidently blocks any understanding of (13) as a complete question.

- (13a) Was {[a vet][saw [the [dog [that ___ found]_{RC}]]]}
 |_____*_____|

Likewise, the violation in (12a) is severe enough to keep (12) from being ambiguous.

It is worth stressing that the (constrained) transformational notion of natural language grammar does not bear just on the relation between word-strings and interpretations. It also bears on the relation between word-strings and pronunciation. English allows contractions like those in (14-16).

- (14) Who do you *want to* kiss ==> Who do you *wanna* kiss
 (15) What do you think *is up* there ==> What do you think *'s-up* there
 (16) What do you think *it is doing* up there ==> What do you think *it's-doing* up there

perceptual, signalling, and conceptual capacities—provide the original (and perhaps best) arguments for the existence of a nontrivial language faculty; see Chomsky (1965, etc.).

Of course, not all theoretically interesting limitations on ambiguity are due to constraints on displacement of lexical items. And even given a proposal about how a word-string can(not) be naturally structured, there are often theoretically interesting questions to ask about why a given structured word-string fails to have the meanings it fails to have. While (19) has the readings indicated in (19a) and (19b), it does not have the reading indicated in (19c).

(19) The senator called the millionaire from Texas

(19a) The senator called the millionaire, and the millionaire was from Texas

(19b) The senator called the millionaire, and the call was from Texas

(19c) #The senator called the millionaire, and the senator was from Texas

One can hypothesize, plausibly, that (19) is structurally ambiguous as between (19 α) and (19 β);

(19 α) {{The senator} [called [the millionaire [from Texas]]]}

(19 β) {{The senator} [[called [the millionaire]] [from Texas]]}

where the former is a sentence with ‘the millionaire from Texas’ as the direct object of ‘called’, while in the latter, ‘from Texas’ is an adjunctive phrase modifying ‘called the millionaire’. And let’s assume that the word-string in (19) cannot be structured so that ‘The senator’ and ‘from Texas’ form a constituent. This still leaves the question of why (19 β) fails to have the meaning expressed with (19c).

One can imagine a language in which (19 β) is used to say that the senator satisfied two conditions: he called the millionaire, and he was from Texas. The phrase [millionaire [from Texas]] can actually be used as a predicate that imposes a conjunctive condition on individuals. So *why* can't the phrase [[called the millionaire][from Texas]] also be used this way? One can say that structured word-strings (as opposed to signals) are unambiguous, and that (19 β) has the meaning indicated in (19b). But this is the fact to be explained. As a start, one might well follow Davidson (1967) in representing the meaning in question—reported colloquially with (19b)—as in (19b').

(19b') $\exists e\{\text{The}(x):\text{senator}(x)\{\text{the}(y):\text{millionaire}(y)[\text{Called}(e, x, y) \ \& \ \text{From}(e, \text{Texas})]\}\}$

Then one can try to explain why ‘from Texas’ must be understood as a predicate linked to the ‘e’-variable, as opposed to the ‘x’-variable. Alternatively, one could formalize (19b) as follows:

$\exists e\{\text{The}(x):\text{senator}(x)[\text{Agent}(e, x)] \ \& \ \text{Past-Calling}(e) \ \& \ \text{the}(y):\text{millionaire}(y)[\text{Theme}(e, y)]\}$; where the hypothesis would be that *every* word in (19) corresponds to a predicate linked to the ‘e’-variable, perhaps via certain thematic relations associated with certain grammatical relations (see Parsons [1990], Schein [1993, forthcoming], Pietroski [in press]). But in any case, it is hard to see how the 2-but-not-3-way

ambiguity of (19) can be accounted for without some such eventish hypothesis about how grammatically structured word-strings are understood. Similar remarks apply to the prepositional phrase in (20),

(20) Nora heard Fido bark in her apartment

which can be understood as a modifier of either verb; see Higginbotham (1983), Vlach (1983).

Constraints on anaphoric pronouns have been much discussed. So I'll just quickly remind that in (21), 'himself' cannot be understood as referentially dependent on 'Pat'; while in (22), it must.

(21) Pat said he thinks Chris should wash himself

(22) Pat wants to feed Chris and wash himself.

The standard explanation is that in each case, 'himself' must have a "local" antecedent. But in (22), the complement of 'wants' has a covert subject that must be understood as coreferential with 'Pat':

Pat_i wants [PRO_i to [feed Chris and wash himself_i]]; see Chomsky (1981, 1986).

A related point is that the famous (23) is roughly synonymous with (23a), not (23b).

(23) John is eager to please

(23a) *John* is eager that *he* please relevant parties

(23b) *John* is eager that relevant parties please *him*

By contrast, (24) is roughly synonymous with (24b), not (24a).

(24) John is easy to please

(24a) It is easy for John to please relevant parties

(24b) It is easy for relevant parties to please John

So there is evidently an interesting constraint on the *interaction* between lexical meanings and more general facts about how the unpronounced arguments of 'please' are understood. Or put another way, there is a constraint on the kinds of meanings that lexical items like 'eager' and 'easy' can have. One can imagine a language with words 'eeger' and 'eezy', such that word-strings homophonous with (23) and (24) are ambiguous—or a language in which the analog of (23) has only the meaning given with (23b), and the analog of (24) has only the meaning given with (24a). But for whatever reason, the human language faculty does not let us lexicalize concepts of eagerness and easiness in these ways.


Implicit in these explanations is the (independently plausible) assumption that certain grammatical relations have specific interpretative significance. The most obvious illustrations are the facts that suggest *thematic* constraints on how the arguments of action verbs can be related to kinds of event participants. Given a transformational grammar, one needs to explain why (25), structured as in (25 α), has no reading according to which there was a stabbing of Brutus by Caesar.

(25) Brutus stabbed Caesar

(25 α) {Brutus [stabbed Caesar]}

And even given constraints that exclude the logically possible transformation shown in (25β),

(25β) {Brutus {stabbed {Caesar { ___ [___]}}}}



explanatory work remains. We want a theory that rules out the following potential interpretation of (25α): $\exists e[\text{Theme}(e, \text{Brutus}) \ \& \ \text{Past-Calling}(e) \ \& \ \text{Agent}(e, \text{Caesar})]$. Since no natural language associates signals with meanings in this way, the obvious hypothesis is that the grammatical relations *being the external argument of* and *being the internal argument of* are each constrained to have a certain kind of significance, at least in sentences like (25) understood in terms of quantification over events; see...Baker (1988, 1997)...Pesetsky (199x).

As noted in discussions of “conservativity” (see Barwise and Cooper [1981]), an analogous point applies to determiners—quantificational expressions like ‘every’ and ‘no’ that can combine with a noun to form a phrase. Sentence (26) is synonymous with (26a).

(26) Every cow is brown (26a) Every cow is a cow that is brown

Likewise, (27) is synonymous with (27a),

(27) Some cows are brown (27a) Some cows are cows that are brown

and similarly for replacement of ‘Some’ with ‘No’, ‘Most’, ‘Three’, and even complex expressions like ‘More than seventeen but fewer than a million’. If we think of determiners as corresponding to relations between predicates—e.g., the (smallest) relation that a predicate *C* bears to a predicate *B* iff every satisfier of *C* is a satisfier of *B*—then determiners correspond to relations with the following property: if *C* bears the relation to *B*, then *C* bears the relation to the complex predicate formed by conjoining *C* with *B*. And many relations between predicates are *not* relations that have this property. Natural languages somehow enforce a substantive constraint on the interpretation of expressions like (26). If the sound of ‘every’ is understood as a determiner, it cannot be understood as signifying (for example) the relation that *C* bears to *B* iff the satisfiers of *C* are equinumerous with the satisfiers of *B*.

Of course, the concept of equinumerosity can be lexicalized, just not with a determiner. The concept of universal quantification can be lexicalized with a determiner—and in other ways. (Think of sentential adverbs or the focus operator ‘only’.) But there is no natural language with a determiner ‘Ryev’ such that [[Ryev cow] is brown] would mean that every brown thing is a cow, or equivalently, that everything is a cow if brown. So while the concept of universal quantification can be lexicalized with a determiner, natural language apparently disallows the logically possible word ‘Ryev’ as a

determiner; see Pietroski (2004, in press) for a proposal. And this is hardly the only constraint that natural language imposes on “logical” vocabulary.

A large body of recent work suggests that if a word like ‘or’ is used to signify disjunction (as opposed to conjunction, or cats), it must be understood as a sign for inclusive as opposed to exclusive disjunction; see Crain et.al. (forthcoming) and references there. Despite the fact that many uses of ‘or’ invite a contrary view, young children who fully grasp the conversational import of (28) also know what (29) means. And they know that (30) requires that every cat be brown, while (31) does not.

- (28) You may have cake or a cookie (29) Don’t kick the dog or pull its tail
(30) Every dog or cat is brown (31) Every cat is brown or black

Or to take another kind of example, sentence (32) fails to have the reading indicated in (32b)

- (32) It is false that every senator lied
(32a) $\neg\forall x:\text{Senator}(x)[\text{Lied}(x)]$
(32b) $\forall x:\text{Senator}(x)[\neg\text{Lied}(x)]$

Likewise, (33) fails to have the reading gestured at paratactically in (33b).

- (33) George rejected the claim that every senator lied
(33a) Rejected{George, that claim $\implies \forall x:\text{Senator}(x)[\text{Lied}(x)]$ }
(33b) $\forall x:\text{Senator}(x)\{\text{Rejected}\{\text{George, that claim } \implies \text{Lied}(x)\}\}$

And (34) fails to have the reading reported with (34b).

- (34) Most dogs chased few cats
(34a) Most dogs are such that they chased few cats
(34b) Few cats are such that most dogs chased them

All of which suggests that quantificational expressions are displaced, covertly in English, *but only so far*.

Before concluding this section, let me also offer a reminder that some structural notion like c-command is sure to play an important role in any plausible account of how word-strings are related to meanings. And this is not “just” because we need such a notion to characterize constraints on antecedence and displacement. Note that in (35) but not (36), ‘or’ is a constituent of a phrase that combines with ‘never’ to form a phrase.

- (35) The boy who cried never expected to see tigers or bears at the party
(36) The boy who never cried expected to see tigers or bears at the party

And this otherwise minimal difference has dramatic interpretive effects.

By themselves, such facts do not establish that the language faculty is a distinctively human innately specified cognitive system specific to language—as opposed to a general learning device, with

I-languages products of this device in response to linguistic experience. As stressed in the recent empiricist revival, there can be constraints that reflect (in surprising ways) the nature and operation of devices designed to detect statistical regularities in the environment and then behave in accordance with those regularities; see Elman (200x), etc. But this point has little force, absent accounts of the *specific* constraints that linguists discover. And examples like (1-36) are “opening acts” in more elaborate and detailed nativist arguments, which often involve convergence of many independent lines of research concerning both adults and young children, including cases of children acquiring I-languages respecting subtle constraints on constructions that simply do not exist for local adults (but do exist, subject to those constraints, for faraway adults); see for example, Hornstein and Lightfoot (198x), Crain and Pietroski (200x, 200x, 200x), Crain et.al. (forthcoming). For present purposes, though, I take it as a premise that the best poverty of stimulus arguments remain unrebutted—and that the farther linguists go beyond the opening acts, the greater the gap between (i) experience of the sort that every normal child encounters and makes use of for purposes of acquiring an I-language, and (ii) linguistic constraints of the sort respected all normal children.

3. Some Implications

I turn now to the *raison d’etre* of the preceding review: evaluating conceptions of meaning in light of some discoveries and explanations that have emerged from the last fifty years of linguistics. I begin with some views that might seem obviously wrong. But even those who would reject these views at the outset can find variants tempting when the going gets rough. (For truth-conditionalists who continue to find everything said here painfully obvious: my claim will be that given your otherwise excellent views, you shouldn’t be a truth-conditionalist.)

3.1 Deflation, Transduction, and Radical Translation: Insufficient Ambition

It should have been clear long ago that a theory of meaning for a natural language L would have to be more than a mere algorithm that associates each signal of L with its meaning(s) in L. Chomsky (1957, 1965) talked about negative facts. Davidson (1967) talked about events. The idea that theories of meaning should be theories of understanding was “in the air,” well before condensing in Oxford; and facts about how speakers don’t understand natural language are surely germane to hypotheses about how they do. Montague (1970, 1974) showed how to systematically associate natural sentences containing quantificational noun-phrases (like ‘every dog that Pat saw’) with interpretively appropriate sentences of the lambda-calculus. And this achievement, which together with Davidson’s made a more general algorithm seem possible, helped sharpen questions concerning why expressions don’t mean what they don’t mean. For even prior to discoveries that suggested a level of *grammatical* structure (LF), at which

the restricted quantifiers of natural language would be displaced, one could ask why natural sentences fail to have certain meanings specifiable with perfectly fine sentences of the lambda-calculus.⁵ Nonetheless, 21st century philosophy of language still harbors the idea that a theory of meaning is adequate, so long as it provides a finite description of the (endlessly many) facts concerning what expressions do mean.

Horwich (1997, 1998) begins with a plausible idea: the meaning of a complex expression E is a property of E that can be viewed as the result of combining the meanings of E's constituents in ways corresponding to ways (exhibited by E) of combining expressions. And he correctly notes that given this, accounting for the mere compositionality of semantic properties is not hard. In particular, one doesn't need a substantive theory of truth to say that the meaning of E is determined by (i) the meanings of E's basic constituents, and (ii) the relevant ways of combining those lexical meanings. But Horwich also adopts a version of semantic "deflationism" according to which "Understanding one of one's own complex expressions (nonidiomatically) is, by definition, nothing over and above understanding its parts and knowing how they are combined (1997, p. 504)." On this view, if "one has worked out how a certain sentence is constructed from primitive syntactic elements," then "provided one knows the meanings of those elements" one understands the sentence "automatically and without further ado...no further process needs to be involved, leading from these initial conditions to the state of understanding the sentence." Horwich concludes that given the grammatical structure of a complex expression, all a "theory" of meaning/understanding must provide is a specification of what the lexical items mean; where such a specification can be given disquotationally, using "axioms" like '*barked* means BARKED' to report that a certain word means what it does. But this seems wrong. To understand a complex expression E, one must also know how the form of E contributes to the meaning of E; and this imposes substantive constraints, not captured by disquotation, on what lexical items can(not) mean.

For example, one doesn't explain or even describe the relevant facts about (19 β)

(19 β) {[The senator] [[called [the millionaire]] [from Texas]]}

by saying that this structured word-string means what it does because: each word means what it does; and each mode of grammatical combination has the significance it has. We want to know what *kinds* of significance the various lexical items and grammatical relations have, and what *specific* significance the various modes of grammatical combination have—in each case, contrasting the actual facts about natural languages with other coherent possibilities. Likewise, the axiom '*eager* means EAGER' will not help explain why 'John is eager to please' differs from the possible (but unnatural) sentence 'John is eeger to please', meaning that John is eager for relevant parties to please him; see Pietroski (200x, forthcoming).

In *this* respect, Fodor's (199x, 199x) account of language comprehension is little better. Fodor thinks that sentences of spoken languages don't themselves have meanings, and that they don't need any, given that speakers have a language of thought. His line of reasoning is roughly as follows: (1) meanings are compositional; but (2) there are no compositional semantic theories for spoken languages, since (2a) such theories would have to be compositional theories of *truth*—or some such symbol-world relation—and (2b) there are no such theories, as Wittgenstein/Austin/Strawson/Chomsky argue; but (3) there can be theories of truth for Mentalese. On this view, we can and should say that utterances of English sentences mean what they do by virtue of being associated with suitable tokenings of Mentalese sentences. Of course, the relevant notion of association can't be an ordinary notion of *translation* if English sentences don't themselves have meanings. As Fodor notes, the relation has to be more brutally causal, and in this sense more like transduction; but see Matthews (2003) for critical discussion. My own view is that this smells like associationism, that (3) remains a leap of faith, and that (2a) is dubious given (1) and (2b). But whatever one thinks about this, merely associating English sentences with Mentalese sentences doesn't explain why the former *aren't* associated with *other* sentences of Mentalese; see Higginbotham (199x).

One wants to know, for example, why each utterance of (8) fails to have as one of its Mentalese “correlates” the yes/no Mentalese question corresponding to (6).

(6) The child who was fed the soup fed the kitten

(8) Was the child who fed the soup fed the kitten

One leaves the theoretically interesting facts unexplained and unmentioned if one simply supplements a Montague-style algorithm with the hypothesis that the lambda-calculus (or whatever) is the language of thought. And the problem (here) is not that one would owe a theory that associates expressions of Mentalese with TRS conditions; cf. Lewis (197x). Even given a truth-conditional semantics for Mentalese—and some independently characterized facts for such a theory to explain—facts of the sort discussed in section two would remain mysterious. And it won't be enough to posit a language faculty that (for whatever reason) merely generates syntactic structures (distinct from those already available in thought), unless “mere” syntactic structures come with the right constraints on how they can be interpreted. But that said, my point is not that language-of-thought hypotheses are irrelevant.

Just as one can think of the language faculty as having a “phonetic side” that generates instructions for articulatory systems (and coverting signals from perceptual systems), one can also think of it as the language faculty as having a “semantic side” that generates instructions for the creation of complex mental representations out of lexicalizable concepts (and linguistically encoding mental

representations from diverse cognitive domains). This old idea lies near the hear of Chomsky’s (199x, 200x) minimalist program. And I think this conception of meanings—as instructions (generated by the language faculty) for how to compose complex concepts—may be the best currently available option; see Pietroski (forthcoming), drawing on Spelke (200x) and Carruthers (200x). But consider again the point that in natural language, the concepts of eagerness and easiness are lexicalized in specific ways that interact, in constrained ways, with whatever composition principles are relevant to determining what ‘John is eager/easy to please’ does and does not mean. Even if it is true that ‘easy’ is the English label for the concept EASY, this leaves the interesting facts about ‘easy to please’ undescribed and unexplained. Similarly, one wants to know why the concept EQUINUMEROSITY is lexicalizable, but not as a determiner. And so on.

This highlights the need for caution when speaking about the (alleged) “syntax-semantics interface.” Especially given the history of these issues, and the way theories of meaning are standardly formulated, it can be tempting to suppose that word-strings get their meanings in two stages. With regard to a string like ‘The senator called the millionaire from Texas’, *perhaps* processes of “pure syntax” impose (two) structures on the words, without adding anything corresponding to the event-variable in Davidsonian meaning specifications. And *perhaps* the outputs of such processes become the inputs to other distinctive processes that associate grammatically structured word-strings with (representations of) meanings. But absent detailed proposals that explain the relevant negative facts, as opposed to merely encoding them in a preferred formalism, there is little reason for adopting this tendentious conception of *how* the language faculty works its magic. Another possibility is that the hypotheses we (currently) take to be proposals about natural language *syntax* are simply partial hypotheses about how the language faculty associates signals with meanings—and likewise for the hypotheses we (currently) take to be proposals about *semantics*. For all we know, which in these respects is very little, talk of an interface may be more distorting than helpful if it encourages pictures of interpretationless tree-diagrams being handed over for semantic enchantment. It may be that appropriately sophisticated refinements of (19β) and (19b’)

(19β) {[The senator] [[called [the millionaire]] [from Texas]]}

(19b’) $\exists e\{\text{The}(x):\text{senator}(x)\{\text{the}(y):\text{millionaire}(y)[\text{Called}(e, x, y) \ \& \ \text{From}(e, \text{Texas})]\}\}$

will reflect aspects of how the language faculty operates without reflecting interfacing submodules.⁶

We have excellent reason for thinking that syntax as currently conceived plays a crucial role in explaining why and how the language faculty does and doesn’t do what it does and doesn’t do. We also have excellent reason for thinking that this role is not exhaustive, even setting aside phonology (and whatever else is required to “parse” perceptible signals into word-strings). But we shouldn’t assume that

the language faculty minus phonology and syntax will be interestingly like a stipulated semantics for the invented syntax of a formal language. Some such assumption might be plausible, as a working hypothesis, *if* we had independent reason for expecting principles that determine the “well formed formulae” of natural languages. But we don’t—as Chomsky, Higginbotham, and others have stressed. We can pretend otherwise for certain heuristic purposes. But why think that in each I-language, (12)

(12) Was the hiker who lost whistled

is either a wff that feels defective, or a non-wff that gets an interpretation? Similarly for (37-38).

(37) Colorless green ideas sleep furiously

(38) The child seems sleeping

While (39) sounds terrible to me, it sounds fine to my friend Dan (from Ohio).

(39) The lawn needs mowed

Is it a wff for him and not for me? One hardly needs this hypothesis to say that our language faculties have stabilized in slightly different states, with the result that for me but not for him: (39) is not assigned any sentential meaning, at least not without a great deal of effort; whereas (38) sounds somehow wrong to both of us, but not in a way that blocks our assigning the same (obvious) interpretation.

Returning to the main thread, once we reject the idea that a theory of meaning for a natural language L can be an “unambitious” algorithm that merely associates expressions of L with (perhaps mental) representations of what those expressions do signify, it becomes clear that various ways of supplementing such an algorithm won’t help. In particular, it won’t help to add that a given unambitious algorithm is the one that a “radical interpreter” (in Davidson’s [198x] sense) would come up with in attempting to interpret the linguistic behavior of relevant speakers. And for reasons that will become clear, I want to end this subsection by elaborating on the irrelevance of radical interpreters; see Pietroski (forthcoming), drawing Chomsky (199x, 200x).

The radical interpreter—like the Quinean field linguist—represents an arbitrary and unstable mix of two perspectives on linguistic meaning: that of a *scientist*, who will assume as little as possible about the space of possible interpretations, whatever they turn out to be; and that of a *child* who needs to figure out, on the basis of limited evidence, which I-language(s) will work best for purposes of communication. (The idea is that the semantic properties of sentences are somehow constructable from, or at least determined by, facts concerning how utterances of those sentences *could be assigned interpretations* by a rational but otherwise linguistically agnostic creature; so one imagines such a creature trying to figure out, on the basis of limited evidence, what speakers are saying.) But actual scientists do not refuse to consider data unavailable to children. And actual children impose substantive constraints on the space of

possible interpretations for signals. Indeed, one way of summarizing the conclusion of nativist arguments in this domain is that human children are not radical interpreters; see Chomsky (198x, 199x). One can hypothesize that the “heuristics” used by children are dispensable in favor of a more intensive search of the preferred evidence. But this assumption is quite implausible. Radical interpreters won’t figure out that expressions of natural language are *unambiguous* in the ways they are.

Correspondingly, the semantic facts about natural languages are not determined by what radical interpreters would conclude. One might be led to think otherwise by the following chain of reasoning: theories of meaning are theories of truth; so we need a conception of meaning/understanding according to which theories of truth can be theories of understanding; and any such conception will treat meaning, in effect, as radical interpretability. But if the last premise is correct—and many have endorsed it—we have an argument against the first. This is not to deny that the environment is responsible for certain aspects of intersubjective stability with regard to how we *use* linguistic expressions to talk about things. (Though such stability may presuppose rather than explain linguistic meaning.) The important point here is that if humans share a biology/psychology that imposes substantive constraints on how linguistic signals can be associated with meanings, one will not discover this simply by attending to evidence available to any normal child. Correlatively, a radical interpreter will be forced to look *in the wrong place* for semantic regularities that are due to constraints imposed by human biology/psychology. In his search for sources of intersubjective linguistic stability, especially with regard to how speakers use expressions to make true claims, a radical interpreter will be forced to exaggerate the relevance of the shared environment. So while radical interpreters will tend to look for associations between words and “things referred to,” this is no reason for thinking that theories of meaning/understanding should make such associations.

For similar reasons, there is little point in supplementing an unambitious algorithm (that maps signals of a language to their meanings) with the claim that it associates words with “things” in the relevant environment. Given arguments for a nativist conception of the language faculty, we cannot just assume that the main task in natural language semantics is to provide theories that associate signals with aspects of the environment—as opposed to explaining why I-languages do and don’t associate signals with meanings (whatever these turn out to be) in specific ways. If investigation suggests that negative facts are more theoretically interesting than (alleged) facts about signal-environment relations, we should take this as a cue.

3.2 The Externalist Language Faculty: Insufficiently Autonomous

One might well take everything said so far to suggest that the best way to defend truth-conditional semantics is the way suggested by Higginbotham (1985) and elaborated in detail by many: combine a

Chomskyan conception of the language faculty with the idea that theories of meaning are theories of understanding, and hypothesize that theories of truth provide the best overall conception of the what I-languages associate signals with. This supplements Davidson's original bold conjecture with a substantive hypothesis about the human language faculty, as opposed to supplementing it with claims about verbal behavior and Radical Intepreters.⁷ The picture is a compelling one, not least because it promises to synthesize the great ideas from the 60's and early 70s. It can seem that Chomsky offers what is needed to account for the negative facts, while Davidson offers what is needed to treat a theory of the language faculty as a theory of understanding, as opposed to a mere theory of linguistic forms. And I readily grant—indeed, I'd argue at length—that pursuing this idea has led to good things. But so far as I can tell, the real virtues of truth-conditional semantics pursued in this way have far less to do with truth (reference, satisfaction) than appearances suggest.

Correspondingly, I think it is very important to distinguish the following two claims: (a) we can combine our best linguistic theories, which incorporate insights from the Davidsonian program, with truth-theoretic models of I-languages; (b) I-languages associate linguistic signals with (functions from contexts to) TRS conditions, and this is what our best linguistic theories say (correctly understood). The first remark may be true. But it amounts to little more than the claim that whatever meanings are—i.e., whatever I-languages associate signals with—theorists can stipulate certain interpretations for the formalism we use to talk about the syntactic structures of natural language. And while this may not be entirely trivial, I assume that (b) is the thesis of interest to truth-conditionalists (like my former self).

At this point, after so much setup, I can state the main poin briefly: this is a *very* tendentious conception of a human faculty that creates expressions from lexical items in accordance with its own (autonomous) constraints. In many ways, the language faculty seems to be a cognitive system that doesn't give a whit about truth (or communication). And while theories of truth for natural languages might be desirable, in various respects, this doesn't give us reason for thinking that stable states of the language faculty associate linguistic signals with TRS conditions. It sounds good to say that truth is the central notion in a theory of meaning—until one takes natural natural languages to be I-languages, which owe much of their character to innately specified aspects of the language faculty. For it seems overly optimistic to suppose that *truth* is a central notion in theories of a largely innate aspect of human psychology. (This makes it seem as though rationalism does, after all, require a beneficent creator.) More importantly in my view, if we think about the kinds of phenomena discussed in section two, talk of TRS conditions adds nothing to the explanation of such facts; and in some cases, such talk impedes theoretical understanding. In the remaining pages, I'll spell out this last remark a bit, though not yet enough.

As I keep noting, it's hard to see how we can even start accounting for the interesting facts concerning (19) without a Davidsonian event analysis of some kind.

(19) The senator called the millionaire from Texas

And historically, such examples were developed in the context of attempts to specify truth theories of natural languages; for initially, it seemed that adverbial constructions presented an immediate stumbling block. But as further study revealed, it is unclear just what event-variables in *semantic* theories range over. And it seems especially unlikely that they range over things in terms of which a theory of *truth* for English would specify a truth-condition for (19). It turns out that for purposes of providing theories of understanding, we need a notion of 'event' that lets one draw distinctions that seem to be metaphysically otiose. Higginbotham offers a now well known example. If a drinker downs his pint (continuously) over thirty seconds, there was an event of drinking a pint of beer in thirty seconds, and there was an event of drinking beer for thirty seconds. One might have been inclined to say, in Davidsonian fashion, that here we have one event and two descriptions. But then given event analyses, it would follow that there was an event of drinking beer *in* thirty seconds; yet in natural language, *'He drank beer in thirty seconds' (see Tenny [199x]). This suggests that natural language cares about the difference between drinking beer, in a case where the quantity drunk was a pint, and drinking a pint of beer. And *prima facie*, this is distinction without difference so far as truth is concerned. Schein (1993) argues that semanticists are driven to distinguish the facing of Carnegie Hall by Carnegie Deli from the facing of Carnegie Deli by Carnegie Hall, the preceding of 3 by 2 from the preceding of 2 by 3, and so on; see also Pietroski (1998).

With heroic measures, one can introduce the structure needed to account for how speakers understand the relevant sentences, but then ignore unwanted aspects of that structure when it comes to preserving the idea that theories of meaning are theories of truth; see Schein (2002) for illuminating discussion. Though if heroic measures are required at every turn, that may be telling. And as soon as we turn to thinking about how grammatical relations are related to thematic roles, examples like (37)

(37) The rock broke the window

suggest that rocks can be Agents in the sense of 'Agent' that matters for theories of meaning; see Baker (1997) for discussion and defense. While this need not be a problem with regard to constructing theories of the language faculty, one wonders whether an honest theory of *truth* for English can really assign a truth-condition to (37) in terms of the rock being an Agent.

We can, if we like, say that the word 'eager' is satisfied by states of eagerness, while 'easy' is satisfied by easinesses. But if we intend this as part of a serious hypothesis about how the language faculty lexicalizes a certain concept (or family of concepts), we owe an explication of what we mean by

phrases like ‘satisfied by’ and ‘states of eagerness’. Given that Tarski’s technical notion is not obviously applicable to predicates of natural language, almost all of which are vague and sensitive to communicative import, one can hardly assume that the homophonous phrase means what it means in Tarski’s theory; see Pietroski (forthcoming). But set this aside. The more pressing concern here is that to account for the relevant negative facts, we need to say enough to make it clear that English differs from the imagined language with speakers who might grow up to say things like the following: the word ‘eeger’ is satisfied by states of eegerness. In part for this sort of reason, Higginbotham (199x) urges truth-conditionalists to provide “elucidations” of lexical meaning—and not to be hostile to lexical analysis. This strikes me as exactly the right move to make. But until we have a significant number of theoretically interesting elucidations specified in truth-conditional terms, appeals to TRS conditions look like placeholders awaiting the arrival of genuinely explanatory notions.

Once we have an explanation for the conservativity of determiners, it will presumably be possible to encode that explanation in a semantic theory according to which ‘Every cow is brown’ is true iff every cow is a cow that is brown. But is there any reason for thinking that talk of truth-conditions will be essential to the explanation? Likewise, once we discover that English ‘or’ signifies inclusive disjunction, we can encode this in familiar truth-conditional terms; though we could have equally well encoded the contrary hypothesis. But if the arguments I alluded to are on the right track, kids have no choice: in (29) and (30), ‘or’ must be understood so that the sentences are equivalent to (29a) and (30a); while (31) must be understood as inviting the (cancellable) implicature reported in (31a).

- (29) Don’t kick the dog or pull its tail (30) Every dog or cat is brown
 (29a) Don’t kick the dog, don’t pull its tail
 (30a) Every dog is brown, and every cat is brown
 (31) Every cat is brown or black
 (31a) Not every cat is brown, and not every cat is black

One wants to know why the language faculty forces kids to hear language in the ways they must. And it doesn’t seem that appeal to TRS conditions will help. Likewise, one can model the significance of various modes of grammatical combination in terms of TRS conditions. But this will be exactly as useful as the corresponding specifications of lexical meanings in the theoretically interesting cases.

I won’t here discuss constraints on anaphoric pronouns, since this touches too closely the kinds of Wittgenstein/Austin/Strawson/Chomsky-considerations I wanted *not* to rely on here. But in light of examples like (38-40), one might well want a concept of meaning that doesn’t require one to specify the meanings of pronouns in terms of TRS conditions; see Pietroski (2003, forthcoming).

- (38) France is hexagonal, and it is a republic
 (39) The red book is too heavy, though it was favorably reviewed,
 and the blue one is boring, though everyone is reading it
 (40) If you ask the average man’s wife whether he likes round squares,
 she’ll say that he doesn’t, but I think he does like them

More in keeping with the thrust of this paper, let me stress again that constraints on displacement are interesting in part because they are so bizarre from a logical point of view. Why forbid extraction from a relative clause? (Remember: we can’t assume that this is an arbitrary constraint of a syntax whose wffs are the only expressions that get interpretations.) And why should “wh-traces” affect pronunciation?

Given examples like (32),

- (32) It’s false that every senator lied

it looks like the language faculty is a system operating in accordance with its own principles, with regard for the relation between quantification and truth. It generates expressions that have the properties they have. Speakers find ways to use those expressions, occasionally as devices for making truth-evaluable assertions; and meaning constrains use, in subtle and interesting ways. But the more we learn about the deep properties of the language faculty, the less plausible “truthy” conceptions of the faculty become. Or so it seems to me. (Chomsky offers such remarks regularly. I make no claim to originality here.)

To be sure, one can provide an algorithm that associates (3b) with an appropriate truth-condition. And one can appeal to a prohibition against extraction from a relative clause in explaining why (3) is unambiguous. But appeal to TRS conditions does no explanatory work here. One *could* just as well assign a truth-conditional interpretation to (3a).

- (3) Was the hiker who lost kept walking in circles
 (3a) Was {[the [hiker [who __ lost]_{RC}]][kept walking in circles]}
 |_____ * _____|
 (3b) Was {[the [hiker [who lost]_{RC}]][__ kept walking in circles]}
 |_____|

Indeed, if a truth theory for English helps at all in explaining why it is odd to say that hiker (as opposed to a song) was whistled—and here is a place where one might expect elucidations to help—then the truth-conditional interpretation assigned to (12a) will be *less* defective than the one assigned to (12b).

- (12a) Was {[the [hiker [who __ lost]_{RC}]][whistled]}
 |_____ * _____|
 (12b) Was {[the [hiker [who lost]_{RC}]]*[__ whistled]}
 |_____|

Which suggests what we should have expected in the first place. Explanations of the negative facts illustrated with (3a) and (12a) will continue to flow from claims about *how* I-languages associate signals with meanings, not proposals about how to characterize meanings in terms of TRS conditions. On the contrary, such proposals look rather like fifth wheels.

4. Friendly Conclusion

Having put the point that strongly, it's time to end by backing off a bit.

Once we learn to ignore the distractions of radical interpreters, we can ask how much room the language faculty leaves for intersubjective linguistic stability due to our shared environment. And appeals to word-world relations may help us characterize such stabilities in interesting ways. My fairly rabid nativist tendencies, fed by studies of creolization and children blissfully acquiring languages their parents do not speak, lead me to be suspicious. But that is another, longer topic. And I want to end on a note of concession.

Whether or not Davidson's bold conjecture is true, it helped us see that theories of meaning can be formulated in terms of conditions that something must satisfy to *be a (Semantic) Value of an* expression—or more generally, drawing on Boolos (1998), conditions that some things must satisfy to be Values of expressions (with zero or more variables) relative to assignments of Values to variables. This makes it possible to formulate theories of meaning for a language without associating each expression of the language with some thing that is *the meaning* of the expression. This lets us frame theoretically interesting questions about lexical meanings, and how they compose, in terms of the kinds of Values various expressions have; see Higginbotham (1985). This is extremely useful, so long as we don't press the idealization too far and suppose that the semantic typology thus introduced requires association of expressions with TRS conditions.

On the contrary, another virtue of the Davidsonian framework is that it shows us how much can be done, starting with the following very spare assumptions: declarative sentences are associated with a distinctive range of (Semantic) Values, which reflect the ways that sentences can be evaluated; relative to any context, each sentence is associated with exactly one Value from the range of potential Values for sentences; synonymous sentences are always associated with the same Value; one of the potential Values for sentences—call it '**t**'—is such that every tautology (like 'Pat sang if Pat sang') has this Value; and a sentence S1 semantically implies another sentence S2 iff relative to every context, S1 has the Value **t** *only if* S2 has the Value **t**. This requires at least two potential Values for sentences to avoid the absurd consequence that each sentence implies every sentence. But the minimal assumption that there are exactly two—call the other one '**f**'—is surprisingly powerful in terms of the explanations it can support. By

contrast, I claim, very little explanatory work is done by the additional hypothesis that a sentence is *true* relative to a context iff it has the Value **t** relative to that context.

The spare Davidsonian typology also makes it clear that talk of sentences having **t/f**-conditions (relative to contexts) need not—and so without supporting argument should not—be understood in terms of associating sentences with “states of affairs.” All we need to get going on the project of constructing substantive theories of meaning is the idea that a sentence is associated with one of two ways of evaluating the sentence, depending on whether or not a certain condition (specified in terms of the relevant lexical items) is met. We need not think of sentences as representing conditions in order to think of sentences as being conditionally associated with “truth values.” But likewise, we need not think of sentences as being conditionally associated with *truth* values to think of them as being conditionally associated with one of exactly two entities that serve as theoretical representations of positive and negative evaluation.

Finally, Davidson and others showed us how to start formulating theories of understanding against the backdrop of an extensional logic, even though substituting ‘Hesperus’ for ‘Phosphorus’ does not preserve the falsity of sentences like (41).

(41) ‘Hesperus is rising’ means that Phosphorus is rising

Not only does this make daily life as a semanticist easier, it suggests that semantic theories—and correspondingly, the language faculty—might prove in *this* respect to be a bit simpler than initial appearances suggested. But like the other virtues of hypothesized truth theories for natural language, this one seems to have less to do with truth *per se*, and more to do with the kind of formalism that seems best for purposes of characterizing I-languages (without supposing that such languages associate signals with TRS conditions). That is no criticism if such virtues are precisely the ones that we should prize in semantic theories, at least for the foreseeable future.

References (still incomplete, and not properly incorporated into the text)

- Austin, J. 1961: *Philosophical Papers*. Oxford: Oxford University Press.
- 1962: *How to Do Things with Words*. Oxford: Oxford University Press.
- Baker, M. 1988: *Incorporation*. Chicago: University of Chicago Press.
- 1997: Thematic Roles and Grammatical Categories. In L. Haegeman (ed.),
Elements of Grammar, Dordrecht: Kluwer, pp. 73-137.
- Barber, A. (ed.) 2003: *Epistemology of Language*. Oxford: Oxford University Press.
- Benacerraff, P. 1965: What Numbers Could not Be. *Philosophical Review* 74: 47-73.
- Bloom, P. 2000: *How Children Learn the Meanings of Words*. Cambridge, MA: MIT Press.
- Boolos, G. 1975: On Second-Order Logic. *Journal of Philosophy* 72: 509-27.
- 1998: *Logic, Logic, and Logic*. Cambridge, MA: Harvard University Press, 1998.
- Burge, T. 1979: Individualism and the Mental. *Midwest Studies* 4: 73-121.
- 1989: Wherein is Language Social. In George (1989).
- Cappellen, H. and Lepore, E. 2003: Radical and Moderate Pragmatics. In Z. Szabo (ed.),
Semantics and Pragmatics (Oxford: Oxford University Press).
- forthcoming: *Insensitive Semantics*. Cambridge, MA: Blackwell.
- Carruthers, P. 2002: The Cognitive Functions of Language.
Behavioral and Brain Sciences 25:261-316.
- Castañeda, H. 1967: Comments. In Rescher (1967).
- Chomsky, N., 1957: *Syntactic Structures*. The Hague: Mouton.
- 1965: *Aspects of the Theory of Syntax*. Cambridge, MA: MIT Press.
- 1966: *Cartesian Linguistics*. Lanham: University Press of America.
- 1969. Quine's Empirical Assumptions. In D. Davidson and J. Hintikka (eds.),
Words and Objections: Essays on the Work of W.V. Quine (Dordrecht: Reidel.)
- 1970: Remarks on Nominalization. In R. Jacobs and R. Rosenbaum (eds.),
Readings in English Transformational Grammar (Waltham: Ginn).
- 1977: *Essays on Form and Interpretation* (New York: North Holland).
- 1981: *Lectures on Government and Binding*. Dordrecht: Foris.
- 1986: *Knowledge of Language* (New York: Praeger).

- 1993: Explaining Language Use. *Philosophical Topics* 20: 205-231.
- 1995a: Language and Nature. *Mind* 104:1-61.
- 1995b: *The Minimalist Program*. Cambridge, MA: MIT Press.
- 1996: *Powers and Prospects*. Boston: South End Press.
- 2000a: *New Horizons in the Study of Language and Mind*.
Cambridge: Cambridge University Press.
- 2000b: Minimalist Inquiries. In R. Martin, D. Michaels, and J. Uriagereka (eds.), *Step by Step: Essays on Minimalist Syntax in Honor of Howard Lasnik* (Cambridge, MA: MIT Press).
- Collins, J. 2003: Horwich's Schemata Meet *Syntactic Structures*. *Mind* 112: 399-432.
- Crain, S. and Pietroski, P. 2001: Nature, Nurture, and Universal Grammar.
Linguistics and Philosophy 24: 139-86.
- 2002: Why Language Acquisition is a Snap. *The Linguistic Review* 19:163-83.
- Davidson, D. 1967a: Truth and Meaning. *Synthese* 17: 304-23.
- 1967b: The Logical Form of Action Sentences. In Rescher (1967).
- 1984: *Inquiries into Truth and Interpretation*. Oxford: Oxford University Press.
- 1985: Adverbs of Action. In Vermazen, B. and Hintikka, M. (eds.),
Essays on Davidson: Actions and Events (Oxford: Clarendon Press).
- 1989: What is Present to the Mind? In J. Brandl and W. Gombocz (eds.),
The Mind of Donald Davidson (Amsterdam: Rodopi).
- 2001. *Subjective, Intersubjective, Objective*. New York: Clarendon Press.
- Dummett, M. 1975: What is a Theory of Meaning? In S. Guttenplan (ed.), *Mind and Language*.
Oxford: Oxford University Press.
- 1976: What is a Theory of Meaning (II)? In Evans and McDowell (1976).
- 1986:A Nice Derangement of Epitaphs: Some Comments on Davidson and Hacking.
In Lepore (1986).
- Evans, G. 1982: *Varieties of Reference*. Oxford: Oxford University Press.
- Evans, G. and McDowell, J. (eds.) 1986: *Truth and Meaning*. Oxford: Oxford University Press.
- Fodor, J. 1987: *Psychosemantics*. Cambridge, MA: MIT.
- 1998: *Concepts: Where Cognitive Science Went Wrong*. New York: Oxford University Press.

- Fodor, J. and Lepore E. 1992: *Holism, A Shopper's Guide*. Oxford: Blackwell.
- 2002: *The Compositionality Papers*. Oxford: Oxford University Press.
- Frege, G. 1879: *Begriffsschrift*. Halle: Louis Nebert. English translation in J. van Heijenoort (ed.), *From Frege to Gödel: A Source Book in Mathematical Logic, 1879-1931* (Cambridge, MA: Harvard University Press, 1967).
- 1982: Function and Concept. In P. Geach & M. Black, trans., *Translations from the Philosophical Writings of Gottlob Frege* (Oxford: Blackwell, 1980).
- Graff, D. 2000: Shifting Sands. *Philosophical Topics* 28:45-81.
- Grandy, R. 1973: Reference, Meaning, and Belief. *Journal of Philosophy* 70: 439-52.
- Hauser, M., Chomsky, N. and Fitch, W. 2002: The faculty of language. *Science* 298:1569-1579.
- Herburger, E. 2001: *What Counts*. Cambridge, MA: MIT Press.
- Hermer, L. and Spelke, E. 1994: A geometric process for spatial reorientation in young children. *Nature* 370:57-59.
- Hermer, L. and Spelke, E. 1996: Modularity and development: the case of spatial reorientation. *Cognition* 61:195-232.
- Hermer-Vazquez, L., Spelke, E., and Katsnelson, A. 1999: Sources of flexibility in human cognition. *Cognitive Psychology* 39:3-36.
- Higginbotham, J. 1983a: Logical Form, Binding and Nominals. *Linguistic Inquiry* 14:395-420.
- Higginbotham, J. 1983b: The Logical form of Perceptual Reports. *Journal of Philosophy* 80:100-27.
- 1985: On Semantics. *Linguistic Inquiry* 16: 547-93.
- 1986: Linguistic Theory and Davidson's Program, in Lepore (1986).
- 1989a: Knowledge of Reference. In A. George (ed.), *Reflections on Chomsky* (Oxford: Blackwell).
- 1989b: Elucidations of Meaning. *Linguistics and Philosophy* 12: 465-517.
- 1990: Contexts, Models, and Meanings. In R. Kempson (ed.), *Mental Representations: The Interface Between Language and Reality* (Cambridge: Cambridge University Press).
- 1993: Grammatical Form and Logical Form. *Philosophical Perspectives* 7: 173-96.
- 1994: Priorities in the Philosophy of Thought. *Proc. Aristotelian Society* supp. vol, 20: 85-106.

- 1998: On Higher-Order Logic and Natural Language. *Proc. of the British Academy* 95: 1-27.
- Hornstein, N. 1984: *Logic as Grammar*. Cambridge, MA: MIT.
- Hornstein, N. and Lightfoot, D. 1981: *Explanation in Linguistics*. London: Longman.
- Horwich, P. 1997: The Composition of Meanings. *Philosophical Review* 106: 503-32.
- 1998: *Meaning*. Oxford: Oxford University Press.
- Jackendoff, R. 2002: *Foundations of Language*. Oxford: Oxford University Press.
- 1997: *The Architecture of the Language Faculty*. Cambridge, MA: MIT Press.
- 1990: *Semantic Structures*. Cambridge, MA: MIT Press.
- Kaplan, D. 1989: Demonstratives. In J. Almog, J. Perry, and H. Wettstein, eds.,
Themes from Kaplan (New York: Oxford University Press).
- Katz, J. and Fodor, J. 1963: The Structure of a Semantic Theory. *Language* 39: 170-210.
- Kratzer, A. 1996: Severing the External Argument from its Verb. In J. Rooryck & L. Zaring
(eds.), *Phrase Structure and the Lexicon*. Dordrecht (Kluwer Academic Publishers).
- Kripke, S. 1980: *Naming and Necessity*. Cambridge, MA: Harvard.
- ms: Reference and Existence: The John Locke Lectures for 1973.
- Larson, R., Ludlow, P., and den Dikken, M, forthcoming: Intentional Transitives.
- Larson and Segal 1995, *Knowledge of Meaning* (Cambridge, MA: MIT).
- Laurence, S. and Margolis, E. 2001: The Poverty of Stimulus Argument.
British Journal for the Philosophy of Science 52: 217-276.
- Lepore, E. (ed.) 1986: *Truth and Interpretation* (Oxford: Blackwell).
- Lepore, E. 1983: What Model-Theoretic Semantics Cannot Do. *Synthese* 54: 167-87.
- Lewis, D. 1972: General Semantics. In D. Davidson and G. Harman (eds.),
Semantics of Natural Language (Dordrecht: Reidel).
- Matthews, R. 2003: Does Linguistic Competence Require Knowledge of Language?
In Barber (2003).
- May, R. 1985: *Logical Form: Its Structure and Derivation*. Cambridge, MA: MIT Press.
- McGilvray, J. 1998: Meanings Are Syntactically Individuated and Found in the Head.
Mind and Language 13: 225-80.
- McGilvray, J. 1999: *Chomsky: Language, Mind and Politics*. Cambridge: Polity Press.

- Montague, R. 1970. *English as a Formal Language*. Reprinted in Montague (1974).
 —1974: *Formal Philosophy*. New Haven: Yale University Press.
- Moravcsik, J. 1975: *Understanding language*. The Hague: Mouton.
 —1998: *Meaning, Creativity, and the Partial Inscrutability of the Human Mind*. Stanford: CSLI.
- Parsons, T. 1990: *Events in the Semantics of English*. Cambridge, MA: MIT Press.
- Pietroski, P. 1998: *Actions, Adjuncts, and Agency*. *Mind* 107: 73-111.
 —2000a: *Causing Actions* (Oxford: OUP).
 —2000b: *The Undeformed Domain of Semantics*. *Sats, Nordic Journal of Philosophy* 1: 161-76.
 —2002: *Function and Concatenation*. In Preyer and Peters (2002).
 —2003a: *The Character of Natural Language Semantics*. In Barber (2003).
 —2003b: *Small Verbs, Complex Events: Analyticity without Synonymy*.
 In L. Antony and N. Hornstein (eds.), *Chomsky and his Critics* (Cambridge, Blackwell).
 —2003c: *Quantification and Second-Order Monadicity*. *Philosophical Perspectives* 17:259-98.
 —forthcoming: *Meaning Before Truth*. In Preyer and Peters (forthcoming).
 —in press: *Events and Semantic Architecture*. Oxford: Oxford University Press.
- Pustejovsky, J., 1995: *The Generative Lexicon*. Cambridge, MA: MIT Press.
- Putnam, H. 1975: *The Meaning of 'Meaning'*. In K. Gunderson (ed.),
 Language, Mind and Knowledge (University of Minnesota Press).
- Preyer, G. and G. Peters, G. (eds.) 2002: *Logical Form and Language*.
 Oxford: Oxford University Press.
 —forthcoming: *Contextualism in Philosophy*. Oxford: Oxford University Press.
- Quine, W.V.O. 1951: *Two Dogmas of Empiricism*. *Philosophical Review* 60:20-43.
 —1960: *Word and Object*. Cambridge MA: MIT Press.
- Rescher, N. (ed.) 1967: *The Logic of Decision and Action*
 (Pittsburgh: University of Pittsburgh Press).
- Schein, B. 1993: *Events and Plurals*. Cambridge, MA: MIT Press.
 —2002: *Events and the Semantic Content of Thematic Relations*. In Preyer and Peters (2002).
 —forthcoming. *Conjunction Reduction Redux*. Cambridge, MA: MIT Press.
- Schiffer, S. 1994a: *The Language-of-Thought Relation and its Implications*.

- Philosophical Studies* 76: 263-285.
- 1994b: A Paradox of Meaning. *Nous* 28: 279-324.
- 2000: Critical Study: Horwich on Meaning. *Philosophical Studies* 50: 527-536.
- Spelke, E. 2002: Developing knowledge of space: Core systems and new combinations.
 In S. Kosslyn and A. Galaburda (eds.), *Languages of the Brain*
 (Cambridge, MA: Harvard University Press).
- Stainton, R. forthcoming: Meaning and Reference—Some Chomskyan Themes.
 In E. Lepore and B. Smith (eds.), *Handbook of Philosophy of Language*
 (Oxford: Oxford University Press).
- Stanley, J. 2000: Context and Logical Form. *Linguistics and Philosophy* 23: 391-424.
- 2002: Making it Articulated. *Mind and Language* 17: 149-68.
- Stanley, J. and Szabo, Z. 2000: Quantifier Domain Restriction. *Mind and Language* 15:219-62.
- Strawson, P. 1950: On Referring. *Mind* 59: 320-44.
- Tarski, A. 1933: The Concept of Truth in Formalized Languages. Reprinted in Tarski (1983).
- 1983: *Logic, Semantics, Metamathematics* (trans. J.H. Woodger, 2nd edition, ed. J Corcoran).
 Indianapolis: Hackett.
- Taylor, B. 1985: *Modes of Occurrence*. Oxford: Blackwell.
- Travis, C. 1985: On What is Strictly Speaking True.
Canadian Journal of Philosophy 15: 187- 229.
- 1996: Meaning's Role in Truth. *Mind* 105: 451-466.
- Vlach, F. 1983: On Situation Semantics for Perception. *Synthese* 54: 129-52.
- Wittgenstein, G.1953: *Philosophical Investigations*. New York: Macmillan.

Notes

1. References, including Carston, Travis, Recanati
2. I assume that meanings are constant across languages, even if certain meanings cannot be expressed in certain languages.
3. While ‘wanna’-contraction is to the left, ‘is’-contraction is to the right, as (15) illustrates and the Gershwins knew: ’s wonderful, ’s marvelous. Likewise, one can ask if *anyone s’up for pizza*, and greet friends by saying (*what*) *s’up*. By contrast: **What do you think’s up there*. For simplicity, I ignore any intermediate position occupied by a question-word on its way to its final position. Note that in ‘What do you think (that) that is up there’, contraction is not possible; and again, there is an intervening wh-trace. Though evidently, not all covert elements are alike in this respect. In (14), the contraction does cross the subject position of ‘kiss’.
4. A language that allowed for more contraction and displacement would, other things being equal, exhibit more synonymy and homophony. Though to repeat, natural languages show little sign of abhorring (contraction or) ambiguity.
5. And Chomsky (1965) was already arguing for the relevance of transformational grammar to such questions in the context of examples like ‘John is easy/eager to please’. For discussion of the “LF” hypothesis, see May (197x, 1985), Higginbotham and May (1981); see Huang (1995) and Hornstein (199x) for more recent reviews. For discussion of how this hypothesis and Davidson’s event analysis bear on traditional questions about logical form, and how it relates to grammatical form, see Pietroski (2004) and references there.
6. In the early days of vision theory, one *could* have declared that the process of mapping from retinal images to 3-dimensional representations divides into (i) visual syntax—concerning the mapping from the two retinal images to pairs of edge-representations that can be “coaddressed” in subsequent processing, and (ii) visual semantics—everything else. But what would have been the point?
7. See especially Larson and Segal (1995); see also Harman (197x) for an early formulation of similar ideas.. For discussion in a Montagovian idiom, see Heim and Kratzer (1995); though see Pietroski (in press) for arguments that this is not the best idiom, not even with regard to textbook cases.