Character Before Content

Paul M. Pietroski, University of Maryland

Speakers can use sentences to make assertions. Theorists who reflect on this truism often say that sentences have *linguistic meanings*, and that assertions have *propositional contents*. But how are meanings related to contents? Are meanings less dependent on the environment? Are contents more independent of language? These are large questions, which must be understood partly in terms of the phenomena that lead theorists to use words like 'meaning' and 'content', sometimes in nonstandard ways. Opportunities for terminological confusion thus abound when talking about the relations among semantics, pragmatics, and truth. As Stalnaker (2003) stresses, in Quinean fashion, it is hard to separate the task of evaluating hypotheses in these domains from the task of getting clear about what the hypotheses *are*. But after some stage-setting, I suggest that we combine Stalnaker's (1970, 1978, 1984, 1999, 2003) externalist account of content with Chomsky's (1965, 1977, 1993, 2000a) internalist conception of meaning.

On this view, the meaning of a declarative sentence is *not* a function from contexts to contents. Linguistic meanings are intrinsic properties of expressions that constrain without determining truth/reference/satisfaction conditions for expressions relative to contexts. As we shall see, this independently motivated conception of meaning makes it easier to accept the attractive idea that asserted contents are sets of metaphysically possible worlds. This is to reject certain unified pictures of semantics and pragmatics. But we should be unsurprised if some of the facts that linguists and philosophers describe are interaction effects, reflecting *both* sentence meaning *and* asserted content. While it can be tempting to think of semantics as conventionalized (or "fossilized") pragmatics, with meaning somehow analyzed in terms of assertion, I think that meaning is more independent of—and probably a precondition of—assertion and truth. This may

be at odds with some of what Stalnaker says about semantics. But the important points are in the spirit of his work: it isn't obvious *what* the study of meaning (or content, or gold) is the study of; but there are better and worse ways of framing relevant questions.

1. Kripke-Propositions

Let a K-proposition be a set of *ways the world could be*, leaving it open whether and where this notion is theoretically useful. Let A-propositions be whatever the sentential variables in our best logical theories range over. If such variables range over *structured* abstracta with "logical form," as suggested by the study of valid inference since Aristotle, then K-propositions are not A-propositions. But one shouldn't reason as follows: assertions are governed by logic; sets don't have logical form; so asserted contents are not K-propositions. At best, this would be a misleading way of defining 'content' as a word for talking about things with logical form.

If *ways* the world could be are as Lewis (1986) describes them—universes, like the one that includes us, but each with its own spacetime and distinctive inhabitants—so be it. Though it seems less extreme to suppose, with Kripke (1980) and Stalnaker (1976, 1984), that ways *the* world could be are just that: possible states of the one and only universe, which actually includes us, but which might have been configured differently in many respects; where possible states need not be Ludwigian totalities of things. Like Kripke, I am inclined to identify possible world-states with possible histories of *this* universe. But for present purposes, we need not decide what possible states (or configurations) of the world are. And whatever they are, we can use the now standard phrase 'possible worlds' to talk about them, without identifying the actual world—i.e., the *way* the universe is—with the actual totality of things; see Stalnaker (1984). If there is the totality of (actual) things, it would seem to be distinct from the way it is, just as all of the things

seem to be distinct from the ways they are. Though let me stress that the possible worlds, as conceived here, include *all* and only the *logically* possible states of the universe. Put another way, the logically possible worlds are none other than the (metaphysically) possible worlds. We can, if we like, talk about restricted notions of possibility corresponding to sets that include some but not all possible worlds. But the logically possible worlds do not include any ways the universe couldn't really be. So the possible worlds do not exclude any logically possible worlds.[1]

Hence, a K-proposition is a set of logically possible worlds. Though let me enter a caveat, using 'Φ' to abbreviate 'Hesperus (exists and) is distinct from Phosphorus'. I am *not* committed to the following biconditional: it is logically possible that Φ iff there is a possible world at which Φ. On the contrary, there is a possible world at which Φ iff the universe could be such that Φ; and I think Kripke argued persuasively that the universe couldn't be that way. Yet intuitively, it is logically possible that Φ. So I conclude that the complex *sentence* 'It is logically possible that Hesperus is not Phosphorus' does not *mean* that relative to some logically possible state of the universe, the embedded sentence is true (as used). This leaves room for a semantic theory, employing a different technical notion of 'π-World', according to which the complex sentence counts as true iff the embedded sentence is true relative to some π-World.[2] My own preference is for a more Fregean theory that associates each 'that'-clause with a linguistically structured entity whose status as a Semantic Value of 'logically possible' does not depend on its being true relative to some possible state of the universe; see Pietroski (2000). But whatever one thinks about debates concerning what 'that'-clauses and words like 'possible' *mean* in natural language, one can grant that every logically possible world is a way the universe could be without endorsing the following generalization: it is logically possible that ... iff 'that ...' indicates a

nonempty K-proposition.[3] No definition could ensure such a link between logic, 'that'-clauses, and possible states of the universe. These are matters for investigation, not stipulation.

This highlights the question of what K-propositions *are* good for, apart from providing a semantics for invented languages with sentential operators like '□' and '◊', glossed in terms of metaphysical necessity and possibility. It can seem that K-propositions are ill-suited to the study of human communication and thought. Sets of possible worlds are individuated without regard to how speakers/thinkers represent their environment; and two sentences can differ in meaning yet be true (in each context) relative to the same possible worlds. Given plausible assumptions, every property is such that the worlds at which it is instantiated by Hesperus are the worlds at which it is instantiated by Phosphorus. The possible worlds at which two plus two is four are those at which there are infinitely many primes—and these are the worlds at which it is false that Φ. So it can seem like a mistake to characterize meanings or contents in terms of K-propositions.

In my view, this is half right; semanticists need different tools. But as Stalnaker argues, the apparent deficiencies of K-propositions reflect potential virtues that should not be ignored, at least not if representing the universe involves locating the way we take it to be in a space of possibilities. Appealing to K-propositions is a way of talking about possibilities while abstracting away—in so far as such abstraction is possible for creatures like us—from details concerning *how* thinking/speaking creatures like us represent those possibilities. This will be useful when and to the degree that it is useful to distinguish (i) questions about the world represented and the relevant background possibilities from (ii) questions about how it and they are represented.

A set of possible worlds can serve as a mind/language-independent representative of *what* someone is thinking or saying if she thinks or says that Hesperus rises in the evening (or that

Phosphorus rises in the morning). In this technical sense, theorists can treat sets of possible worlds as potential contents: abstract "things" that can be asserted, judged true, doubted, etc. Though to make it explicit that this is a technical sense, we can call K-propositions *kontents*, and then ask whether appealing to them in accounts of human thought and communication is indeed as useful as Stalnaker and others contend. But one cannot object simply by noting, in various ways, that K-propositions play the role they are designed to play as abstractions.

Episodes of asserting that two plus is four differ from episodes of asserting that there are infinitely many primes, presumably in the differing details of *how* the universe gets represented in such episodes. In my view, we shouldn't abstract away from such details when providing theories of *meaning* for natural language, but such abstraction may be perfectly appropriate when talking about *truth*-conditions. Likewise, asserting that Hesperus is Venus differs from asserting that Phosphorus is Venus. But as we'll see in section three, natural language often marks distinctions (relevant to theories of meaning) that are metaphysically otiose. And it may be a *virtue* of appealing to kontents that it leads us to diagnose such distinctions in terms of asserting something in different ways—say, by using sentences with different meanings. (I defer discussion of "two-dimensional" kontents to an appendix.) If the semantic properties of natural language expressions are irredemably human and internalistic, appeal to K-propositions will be ill-suited to certain theoretical tasks in linguistics. Though such appeal may help us articulate and answer various questions concerning the *use* of natural language—especially with regard to heavily world-directed uses, like making truth-evaluable assertions and reporting beliefs. It may be that K-propositions are what we need, and all we're likely to get, for purposes of characterizing a substantive mind/language-independent notion of truth-conditions.

As Stalnaker (2003) emphasizes, philosophers should be especially sensitive to the desirability of a framework that lets us distinguish questions about the truth (or plausibility) of what a person said, from interpretive questions concerning what she said, and questions concerning the meanings of her words. He says that while we "cannot separate semantics from substantive questions before we begin to theorize," philosophers can still try "to separate, in context, questions about how to talk, or about how we in fact talk, from questions about what the world is like (2003, 4-5)." I find much to agree with in these passages where Stalnaker expresses his affinity with Carnap's (1950) project of framing metaphysical questions in terms of a distinction between internal and external questions. But I want to challenge the idea—common to Carnap, Quine, and many others—that semantics is somehow not substantive. Of course, one can define 'semantics' as one likes. But there are substantive constraints on how signals of a human language can be naturally associated with meanings. And this is relevant in the present context.

2. Chomsky-Meanings

In philosophy, as in common parlance, there is a tendancy to equate semantic facts with facts that illustrate Sassurean arbitrariness and the presumably conventional aspects of language use: the French word 'chien', synonymous with 'dog' (and not 'cat'), is typically used to talk about dogs (as opposed to cats); and so on. Similarly, if exactly one of two metaphysicians uses 'proposition' to signify sets of possible worlds, we are apt to say "That's just semantics." Theoretically interesting questions are unlikely to turn on such idiosyncratic facts concerning *which* meaning certain speakers happen to associate with a given perceptible sign.[4] But there is still plenty for semanticists to do, since the interesting facts often concern ways in which natural language *cannot* be understood. Moreover, as discussed by Chomsky and many others, the

constraints seem to reflect a "human language faculty," as opposed to general principles of reasoning, communication, convention, or learning. If this is correct, perhaps semantics should be characterized in terms of this faculty, whose mature states correspond to natural languages.

Extending a famliar slogan: if theories of meaning are theories of understanding, and these turn out be theories of a mental faculty that associates linguistic signals with meanings in constrained ways, then we should try to figure out (in light of the constraints) what this faculty associates signals with.[5] Following Chomsky (2000a), I don't think theories of meaning for natural language will be theories of truth, in large part because I find it implausible that mature states of the language faculty associate signals with truth/reference/satisfaction conditions; see Pietroski (2003a, 2005b). But my aim in this section and the next is just to show that there is motivation and conceptual room for a more internalistic conception of the language faculty, according to which it associates certain signals with instructions for building concepts (much as it associates certain concepts with instructions for generating signals). For these purposes, I assume some familiarity with transformational grammar—and in particular, with the idea that sentences can have unpronounced elements, including traces of displacement operations.

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 less expected thought indicated with (2).

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

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

But (3) has only the meaning indicated with (3b), the yes/no question corresponding to (2).

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

(3a) Was it the case that the hiker who was lost kept walking in circles

(3b) Was it the case that the hiker who lost was kept walking in circles

We hear (3), unambiguously, as synonymous with (3b) *and not* (3a). Likewise, (4)

(4) Was the child who fed some waffles at breakfast fed the kittens at noon

can only be the yes/no question corresponding to the bizarre (6).

(5) The child who was fed some waffles at breakfast fed the kittens at noon

(6) The child who fed some waffles at breakfast was fed the kittens at noon

And we know that natural language is not hostile to ambiguity, given examples like (7).

(7) Solicitors who can duck and hide whenever visiting soliciters may scare them saw every doctor who lost her patience with patients.

If a string of words can be understood as a sentence, but *not* 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 (1967), Travis (1984). One hypothesizes that the transformation indicated in (3β) is licit, while the transformation indicated in (3α) is not.

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

(3β) 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. While (8) and (10) are fine, (9) and (11) are each somehow defective.

(8) The hiker who was lost whistled

(9) Was the hiker who lost whistled

(10) The vet saw each dog that was found (11) Was the vet saw each dog that found

Though (11) is closer to word-salad. We can start to explain this by noting that (9) would be the word-string corresponding to each of the potential transformations indicated in (9α) and (9β);

(9α) Was {[the [hiker [who __ lost]$_{RC}$]][whistled]}
|_____*_____|

(9β) Was {[the [hiker [who lost]$_{RC}$]] ☹[__ whistled]}
|_____|

where '☹' indicates the anomaly of asking whether a *hiker* was whistled—in the way a song can be whistled. This defect is not so severe that it keeps us from hearing (9) as a strange question. (The hiker who lost was whistled? Do you mean that someone whistled to the lost hiker?) But we cannot even hear (11) as a defective way of asking the question indicated with (11α).

(11α) Was {[the vet][saw [each [dog [that __ found]$_{RC}$]]]}
|_____*_____|

Likewise, the violation in (9α) is severe enough to keep (9) from being ambiguous.

Constraints on transformations bear on pronunication as well as interpretation. English allows for contractions like those in (12-14), with blanks indicating wh-traces.

(12) Who do you *want-to/wanna* kiss __

(13) What do you think __ *is-up/'s-up* there

(14) What do you think it *is-doing/'s-doing* __ up there

But in cases like (15-16), contraction is apparently blocked by wh-traces.

(15) Who do you *want* __ *to/*wanna* kiss Chris

(16) What do you think it *is* __ *up/*'s-up* there

One can imagine a language that allows contraction in (16), or one that treats (3) as ambiguous. But in natural language, pronunication is related to interprtetation via grammatical structures that

constrain both, even if the constraints are arbitrary with regard to both. We can easily *produce* the contracted form of (16), just not as a signal with the same (coherent) sentential meaning as the uncontracted form. And we can certainly comprehend the question, expressible with (3a), that cannot be asked with (3). Our inability to impose this sensible interpretation on (3) suggests that natural languages associate word-strings with complex meanings in ways that satisfy constraints on interpretation *independent of* any limitations imposed by (logic and/or) cognitive systems other than the language faculty. Likewise, such associations satisfy constraints independent of those imposed by the cognitive/biological systems responsible for articulation and perception of linguistic signals. Moreover, the relevant constraints on interpretation and pronunciation apparently overlap, suggesting common language-specific factors.

Indeed, evidence that natural languages are governed by "autonomous" constraints—due to the nature of human language, as opposed to perceptual/articulatory/conceptual capacities—is the original (and perhaps best) evidence for a substantive language faculty; see Chomsky (1965, 1986). I offer additional examples in section three. But at this point, we need to sharpen the terminology. Let's say that a *language* is a system for associating signals with interpretations. A *human* language is one that normal human children can acquire in an environment not atypical for members of our species. In so far as 'system' is ambiguous, between a set of abstract rules and a mechanism that instantiates such rules, 'language' is ambiguous. So following Chomsky (1986), let's use 'I-language' to talk about the relevant aspects of human minds/brains: each speaker of a human language has an I-language, by virtue of which she can associate endlessly many linguistic signals—like Japanese *sounds*, or ASL *signs*—with interpretations. This leaves it open what interpretations are. 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. The picture is familiar.

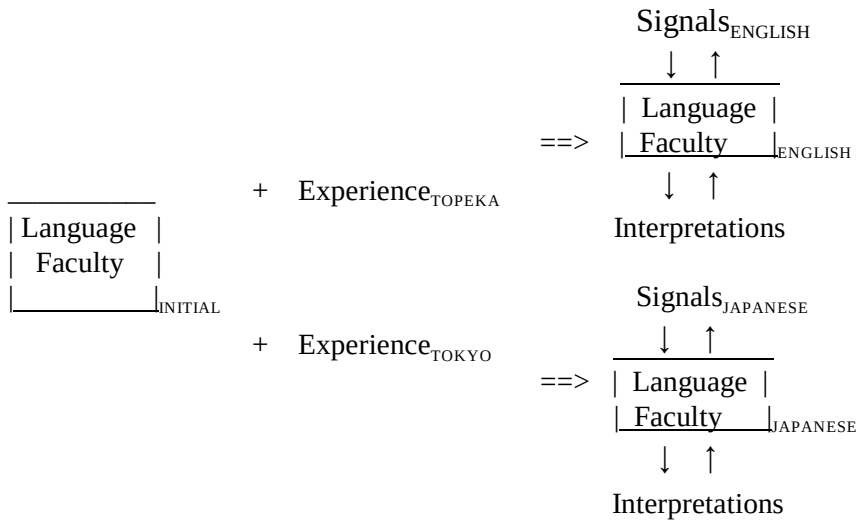

We can go on to ask in which respects, and to what degree, I-languages are: determined by human biology, as opposed to contingent aspects of experience; "encapsulated" from other aspects of human psychology; shared by speakers within a given linguistic community; etc. But the idea is that each "setting" of the language faculty corresponds to the grammar of a possible human language. The range of possible settings is determined by the nature of the human language faculty, which thereby determines which logically possible languages are "options" for a child prior to relevant experience; see, e.g., Chomsky (1981, 1986).[6]

Of course, an I-language may be related to publicly perceptible signals only indirectly. It may even be misleading to talk about the language faculty generating *re*presentations of signals, whose linguistically important properties may not be specifiable independently of the structures generated. Similar remarks apply to interpretations, which remain far less well understood than pronunications. And in any case, it is hardly obvious that I-languages associate signals with

(representations of) truth/reference/satisfaction conditions. One might say instead that each complex expression of a spoken I-language links a pair of instructions for creating a complex sound and a complex concept, leaving it open how (and at what "distance") these instructions are related to entities in the public domain. Indeed, recent trends in linguistics suggest that a grammatical expression may just *be* a pairing (generable by the language faculty) of a "phonological instruction" to articulatory-perceptual systems with a "semantic instruction" to conceptual-intentional systems; see Chomsky (1995, 2000b).

This invites the thought that linguistic meanings are human *Begriffsplans*: blueprints, produced by the language faculty, for constructing concepts from lexicalized elements.[7] On this view, the meaning of 'brown dog that barked loudly' is an instruction for how to build a complex concept out of resources indicated with the constituent words. A theorist might offer a first-pass proposal about certain aspects of this instruction by using notation like the following.

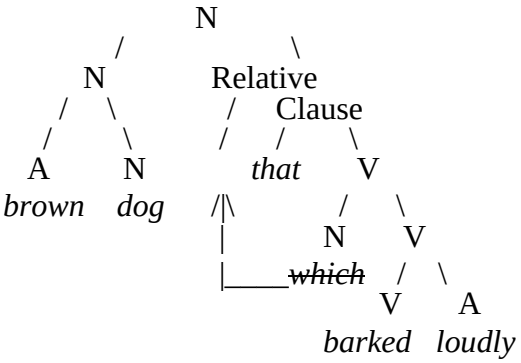

Brown(x) & Dog(x) &

∃e[Agent(e, x) & Past-Barking(e) & Loud(e)]

The logical formalism on the right encodes a hypothesis about how the phrase is understood, as a conjunctive monadic predicate with "thematic" structure. The tree diagram encodes hypothesized grammatical structure, including a displaced covert wh-element, and typology. But the formalism is not self-interpreting. One has to say what symbols like 'x', '∃', 'e', and 'Agent' mean *here*—in the context of theorizing about the human language faculty. And interpreting the variables as

ranging over a doimain of “real world” entities may not be the best way of describing the instructions. For purposes of characterizing *Begriffsplans*, it may be a mistake to abstract away from *how* humans represent the universe. Correlatively, the notation above may reflect two intrinsic aspects of a natural language expression *and not* a proposal about how a mere piece of syntax generated by the language faculty is related to a domain in terms of which one can provide a plausible theory of truth (that can also serve as a theory of meaning) for English.

This leaves ample room for appeal to K-propositions in descriptions of *how humans use* language and concepts to represent the universe. One can ask how ways-we-take-the-world-to-be are located in a space of ways the world could be. But one shouldn’t assume that this is a primary question for the study of linguistic meaning, or that theories of the language faculty should be specified in terms of notions appropriate to investigations based on the kind of abstraction embodied in appeal to K-propositions. This should be obvious, especially in light of Chomsky’s (1986) tripartite distinction of research questions: what do adults know about their languages; how do children acquire such knowledge; and how do humans put such knowledge to use? But the point bears repetition. Theories of *understanding* and theories of *use*, if such there be, may require different kinds of abstraction and idealization. Meaning constrains use, and one shouldn’t expect meaning to be determined by observable facts about how language is used.

Quinean Translators and Davidsonian Interpreters are misguided. They try to figure out what expressions mean, without relying on a language faculty, given only evidence of the sort available to children. But such evidence does not include the (typically negative) facts that support nativist hypotheses about how children (despite their limited experience) figure out what expressions do and don’t mean. The “idealized field linguist” is an empiricist monster—half

scientist, half child—operating without prior assumptions about the constrained character of human languages, but also blind to the evidence that reveals semantically relevant constraints. If any such alien searches for sources of linguistic stability across human speakers, it will be led to exaggerate the role of our shared environment and underestimate the role of our shared biology. Since the alien ignores the most interesting facts, it will conclude that semantics is relatively superficial—a matter of contingent/conventional associations of grammatically structured signals with features of the environment. But if theories of understanding are best viewed as theories of how the *language faculty* associates signals with meanings, as opposed to theories that an alien interpreter would adopt on the basis of limited evidence, perhaps we should take meanings to be human *Begriffsplans* and try to characterize them in terms appropriate to the language faculty.

3. Complex Concepts and Mentalese

The view on offer is, in some respects, like Fodor's (1975, 1978, 1987, 1998) version of the much older idea that linguistic expressions signify mental representations. But there are significant differences. Fodor combines skepticism about truth-conditional semantics for spoken languages with acceptance of the idea that theories of meaning need to be theories of truth, or some such symbol-world relation.[8] Nonetheless, he maintains, there are systematicities that require explanation in terms of a compositional semantics for *some* language that humans use; see Fodor and Lepore (2002). Fodor has independent reasons for positing a language of thought. So he suggests that a spoken language like English has a syntax but no semantics. The idea is that the language faculty associates signals with syntactic structures that must (to be meaningful) be paired, in contexts, with expressions of Mentalese: strictly speaking, a sentence of English does not itself have a linguistic meaning. According to Fodor (1998), sentential utterances are—in

complicated, context sensitive ways that preclude theories of truth for spoken languages—associated with tokenings of Mentalese sentences; where these contextually embedded tokenings do have compositionally determined truth conditions (ceteris paribus, *bien sur*).

From this perspective, expressions of a spoken language are devices for perceptibly indicating concepts whose semantic structure (if any) is independent of the human language faculty. Of course, Fodor doesn't think the relations between syntactic structures and signified concepts are wholly arbitary, as though each spoken sentence were a new word. So he can't really think, despite what he sometimes says, that spoken languages respect no semantic constraints of their own; see Matthews (2003). But in so far as Fodor's view *is* at odds with the idea that spoken sentences have meanings of their own, it makes a mystery of many systematic constraints on possible interpretations of such sentences; see Higginbotham (1994). One wants to know why there is no context in which the Mentalese correlate of (3) is the correlate of (3a).

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

(3a) Was it the case that the hiker who was lost kept walking in circles

(3b) Was it the case that the hiker who lost was kept walking in circles

Why *can't* the string of words in (3) be used to indicate the Mentalese sentence with which we think the thought actually expressed with (3a)? It is not enough to say that (3α) is ungramamtical. For if (3β) is intrinsically *meaningless,* why can't it be used to signify the meaning of (3a), perhaps along with other meanings? (Compare 'bank'.)

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

(3β) Was {[the [hiker [who lost]$_{RC}$]][__ kept walking in circles]}
|_____|

More generally, given a theory of Mentalese, one might be able to specify an algorithm that pairs sentences of Mentalese with Chomskyan syntactic structures in a descriptively adequate way. But one would want still to know why spoken languages are repeatedly interpreted in accordance with this algorithm as opposed to others; cf. Montague (1974). And not all theoretically interesting limitations on ambiguity are due to constraints on displacement of lexical items.

While (17) has readings indicated with (17a) and (17b), it does not have the third reading.

(17) The senator called the millionaire from Texas

(17a) The senator called the millionaire, and the millionaire was from Texas

(17b) The senator called the millionaire, and the call was from Texas

(17c) #The senator called the millionaire, and the senator was from Texas

One can hypothesize, plausibly, that (17) is structurally ambiguous as between (17α) and (17β);

(17α) {[The senator] [called [the millionaire [from Texas]]]}

(17β) {[The senator] [[called [the millionaire]] [from Texas]]}

where (17α) is a sentence with 'the millionaire from Texas' as the direct object of 'called', and in (17β), 'from Texas' is an adjunctive phrase modifying 'called the millionaire'. But even given that (17) cannot be structured so that 'The senator' and 'from Texas' form a phrase, there remains the question of why (17β) fails to have the meaning expressed with (17c). One can imagine a language where (17β) means 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 (17β) has only the meaning/Mentalese-correlate indicated with (17b). But this is the fact to be explained.

As a start, one might follow Davidson (1967b, 1985) in representing the meaning of (17b) as follows: ∃e{The(x):senator(x){the(y):millionaire(y)[Called(e, x, y) & From(e, Texas)]}}. Then one can try to explain why in (17β), 'from Texas' must be understood as a predicate of the 'e'-variable, as opposed to the 'x'-variable. Alternatively, one could formalize (17b) as follows: ∃e{The(x):senator(x)[Agent(e, x)] & Past-Calling(e) & the(y):millionaire(y)[Theme(e, y)]}. Then one can hypothesize that *each* phrase in (17β) must be understood as a predicate of the variable associated with the verb, perhaps via thematic relations associated with certain grammatical relations or prepositions (see Parsons [1990], Schein [1993, forthcoming], Pietroski [2002, 2005a]). But in any case, it is hard to see how the *2-but-not-3*-way ambiguity of (17) can be accounted for without some such eventish constraint on the interpretation of spoken language.

As another example of interesting nonambiguity, note that the famous (18) is roughly synonymous with (18a), not (18b).

(18) John is eager to please

(18a) *John* is eager that *he* please relevant parties

(18b) *John* is eager that relevant parties please *him*

By contrast, (19) is roughly synonymous with (19b), not (19a).

(19) John is easy to please

(19a) It is easy for *John* to please relevant parties

(19b) It is easy for relevant parties to please *John*

So there is evidently a constraint on how lexical meanings interact with more general facts about how the unpronounced arguments of 'please' are understood in (18) and (19). One can imagine a language with the words 'eeger' and 'eezy', such that strings homophonous with (18) and (19)

are ambiguous—or a language in which the homophone of (18) has only the meaning given with (18b), and the homophone of (19) has only the meaning given with (19a). But for whatever reason, the human language faculty does not permit these ways of composing expressions that include lexicalization of notions like eagerness and easiness.

Such constraints tell against the idea, also criticized by Stalnaker (1984), that speakable expressions acquire semantic life only by virtue of being paired with expressions of Mentalese. I do not doubt that we have structured mental representations independent of the language faculty, or that these representations are associated with pronounceable expressions in very complex ways. But prima facie, such associations are constrained by what the spoken expressions mean; the associations don't determine what the expressions mean, except perhaps with regard to genuinely arbitrary aspects of meaning (e.g., that 'chien' is associated with the concept DOG, not CAT). Moroever, if Fodor is right, the human mind exhibits two different kinds of syntax. For absent structural *mismatches* between Mentalese and English sentences, it is hard to see how there could be a compositional semantics for only the former. But one wonders why nature—faced with the problem of getting Mentalese pronounced—would invent a mismatching syntax, thereby saddling children with the task of figuring out how adults associate one kind of syntax with another (as *part* of the task of figuring out what adults are saying).[9]

By contrast, the idea of meanings as *Begriffsplans* fits nicely with the idea that the language faculty allows for a certain kind of integration and expansion of hominid psychology. Various studies suggest that spoken language is implicated in the adult capacity to form—in ways that prelinguistic children and other very clever animals cannot—certain kinds of complex thoughts whose conceptual constituents lie in disparate cognitive domains; see Hermer and

Spelke (1994, 1996), Hermer-Vasquez *et.al.* (1999), Spelke (2002), Carruthers (2002). And whatever else the language faculty lets us do, it lets us *lexicalize* mental representations and *combine* lexical items; where lexicalization involves creating an expression of a certain grammatical type, and the permissible modes of combination depend on grammatical (as opposed to "cognitive") type. This is one of the points Chomsky illustrated with (20),

(20) Colorless green ideas sleep furiously

which is bizarre in many ways, but still a sentence that is not completely incomprehensible. And it would hardly be surprising if hominids without an (activated) language faculty had certain concepts they could not combine, even if such combination would be useful.

Consider the disparate concepts required to understand a phrase like 'every child bigger than the brown dog that did not bark loudly'. One needs a system of "common denominators" to combine these concepts in a coherent fashion—and not just because size concepts are associated with relations between individuals, while the concept of barking is not. We may naturally think about colors as properties *of* surfaces, and think about dogs as things that *have* surfaces, with the result that brute concatenation of the relevant concepts would be unnatural for us. And this kind of mismatch may be common, especially if human cognition is significantly modular. One can think of a *Begriffsplan* as providing, via resources that the language faculty makes available, a way of building complex concepts from elements that otherwise would not fit together.

Of course, the plausibility of this view depends on the plausibility of specific proposals. And one might be suspicious if semantic theories for natural languages *had* to be characterized in terms of the hypothesis that combining expressions typically signifies *function-application*, as in theories inspired by Frege and Montague. For one might doubt that the prelinguistic hominid

mind traffics in such a general and abstract notion. Moreover, if concepts can lexicalized as words whose (Semantic) Values are functions from Values of other lexical items to potential Values of complex expressions, then the idea of meanings as *Begriffsplans* is trivialized. But as the simple examples above already suggest, one can characterize many linguistic meanings compositionally without resorting to descriptively omnipotent notions like function-application. In fact, one can handle all the usual textbook cases and more in terms of predicate-conjunction, existential closure, and a few thematic relations corresponding to certain grammatical relations and prepositions; see Pietroski (2005a), drawing on Boolos (1998), Higginbotham (1985), Parsons (1990), Schein (1993), Larson and Segal (1995). It remains to be seen how much natural language semantics can be done in these terms. But recoding extant work is often easy. And if study of the language faculty suggests that meanings are constrained *Begriffsplans*, we shouldn't immediately conclude that something has gone wrong, on the grounds that some current semantic theories apparently require less restricted conceptual resources. The correct response may be to offer more constrained semantic theories; see Chomsky (1965) for still relevant discussion.

4. Autonomous Constraints, Worldly Truths

Examples like (1-20) do not themselves establish that the language faculty is a largely innate cognitive system specific to human language—as opposed to a general learning device, with I-languages as products of this device in response to linguistic experience. But such examples 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.[10] For present purposes, I take it as a premise that the best poverty of stimulus arguments remain unrebutted, and that more detailed invesigation bolsters the central point: there is a huge gap

between (i) experience of the sort that each normal child makes use of, for purposes of acquiring an I-language, and (ii) linguistic constraints of the sort respected by all normal children. In which case, conceptions of linguistic meaning should be evaluated accordingly: whatever meanings are, children associate them with signals in substantively constrained ways.

If this is correct, as Chomsky has long argued, then it is prima facie implausible that natural language meaning is as tightly related to *truth* as Davidson (1967a, 1984) and others have conjectured.[11] If human I-languages are determined largely by innate aspects of human biology, then absent a benficient deity, it seems unlikely that such languages pair sentential signals with truth conditions. Once we drop the idea that meaning is determined by what Radical Interpreters would say, in favor of studying meaning by investigating the language faculty, why think that the theoretical notions (and idealizations) required to characterize linguistic meanings are apt for purposes of providing theories of truth? Or putting it the other way around, why think the notions appropriate for characterizing the conditions in which assertive utterances are true or false will also be suitable for purposes of characterizing the language faculty? This seems too convenient to be true—even if we think of truth conditions as conditional specifications of truth values, as opposed to propositions or "states of affairs."

One can, of course, hypothesize that truth theories provide the best conception of what I-languages associate signals with; see Higginbotham (1986), Larson and Segal (1995). But we should be clear that this is a bold conjecture about the language faculty. For the empirical virtues of truth-conditional semantics may have little if anything to do with truth *per se*. One can interpret theories that associate sentences with **t** or **f** (relative to assignments of Values to variables) as theories that treat sentences as devices for doing something—evaluable in a binary

fashion—with a concept built in accordance with a certain *Begriffsplan*. Since assertively uttering a sentence is a paradigm of such use, linguistic meanings constrain the truth conditions of assertive utterances. But it does not follow that meanings of expressions, as opposed to asserted contents, should be specified in terms of truth/reference/satisfaction (TRS) conditions.

Correspondingly, it is very important to distinguish 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. While (a) may be true, it amounts to little more than the claim that theorists can stipulate certain interpretations for the formalism we use to talk about the syntactic structures of natural language. And even if such stipulation is useful for certain limited purposes, related to the phenomenon of semantic implication, (b) is a much stronger claim about how the language faculty is (or comes to be) related to the environment. I won't review the arguments that lead me to think (b) is false.[12] But I do want to note one kind of skepticism that is often ignored.

As noted above, it's hard to see how we can even start accounting for the interesting facts concerning (17) without an event analysis of some kind.

(17) The senator called the millionaire from Texas

Historically, event analyses were offered as attempts to specify (partial) truth theories for natural languages; since without them, it seemed that adverbial constructions presented an immediate stumbling block. But as further study revealed, it is unclear what event-variables in *semantic* theories range over; and various considerations suggest that such variables *don't* range over things in terms of which a theory of truth for English would specify a truth-condition for (17). It turns out that for purposes of providing theories of understanding, we need a notion of 'event'

sensitive to distinctions that seem to be metaphysically otiose. Higginbotham offers an example.

If a drinker downs a 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. In eventish formalism: ∃e[Drinking(e) & Theme(e, a pint of beer) & In(e, 30 seconds)]; and ∃e[Drinking(e) & Theme(e, beer) & For(e, 30 seconds)]. One is inclined to say that here, we have one event and two descriptions. But if the 'e'-variable ranges over things that can satisfy multiple descriptions in this way—and one expects the ontology for a theory of truth to be language-independent in this sense—then there was an event of drinking beer *in* thirty seconds: ∃e[Drinking(e) & Theme(e, beer) & In(e, 30 seconds)]. Yet in natural language, 'He drank beer in thirty seconds' is anomolous; see Tenny (1994) for extended discussion. This suggests that natural language cares about the difference between (i) drinking a pint of beer, and (ii) drinking beer, even when the beer drunk amounted to a pint. And this seems to be a distinction without difference so far as truth is concerned.

Schein (1993) argues that semanticists must likewise 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 succeeding of 2 by 3, and so on. With enough ingenuity and tenacity, one can find ways of introducing the structure needed to account for how speakers understand the relevant sentences, and then ignoring 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. But if such measures are required at every turn, perhaps nature is telling us something. And as soon as we turn to thinking about how grammatical relations are related to thematic roles, (21)

(21) The rock broke the window

suggests that rocks can be Agents in the sense of 'Agent' that matters for theories of meaning; see Baker (1997) for discussion and defense; see also Pietroski (1998). While this need not be a problem for internalistic theories of the language faculty—in which 'Agent' can play whatever role it needs to play—one wonders whether an honest theory of *truth* for English can really specify a truth-condition for (21) in terms of the rock being an Agent.

Moreover, and more importantly, it is unclear what appeal to TRS conditions *adds* to explanations of the negative facts that animate study of the language faculty. We can, if we like, say that '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 certain notions, we owe an explication of phrases like 'satisfied by' and 'states of eagerness'. Given that predicates of natural language appear to be vague and sensitive to communicative import, one can hardly assume both that 'satisfied by' means what it means in Tarski's theory, and that claims like *'dog' is satisfied by dogs* are truistic; see Pietroski (forthcoming). But set this aside, along with concerns about the covert indices that will be required by a substantive theory of truth that accommodates all the ways truth can depend on context; see Stanley (2000, 2002), cf. Cappelen and Lepore (forthcoming). The crucial point 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, whose speakers can say things like the following: 'eeger' is satisfied by states of eegerness; and 'John is eeger to please' means that John is eeger to please.

Partly for this reason, Higginbotham (1989) urges truth-conditionalists to provide "elucidations" of lexical meanings. This strikes me as the right move to make, in order to defend a truth-conditionalist conception of the language faculty. But until we have a significant number

of theoretically interesting and confirmed elucidations specified in truth-conditional terms, locutions like ‘satisfied by’ seem to be placeholders for appropriate theoretical notions. And recall that one *could* assign a perfectly coherent truth condition to (9α).

(9α) Was {[the [hiker [who __ lost]$_{RC}$]][whistled]}
|_____*_____|

So *if* a truth theory helps explain why it is odd to say that hiker (as opposed to a song) was whistled—and here is a place where elucidations might well help—then the truth-conditional interpretation that could be assigned to (9α) will be *less* defective than the one assigned to (9β).

(9β) Was {[the [hiker [who lost]$_{RC}$]]*[__ whistled]}
|_____|

All of which suggests what we should have expected: explanations of the negative facts illustrated with (9) flow from claims about *how* I-languages associate signals with meanings, not from proposals about how to characterize meanings in terms of TRS conditions.

In this context, 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?) Natural language quantifiers are also governed by such constraints. For example, (22) has only the reading indicated in (22a),

(22) It is false that every senator lied

(22a) ¬∀x:senator(x)[Lied(x)] (22b) ∀x:senator(x)[¬Lied(x)]

suggesting that ‘every senator’ is displaced but only so far; see Pietroski and Hornstein (2002) for discussion and further references. So it looks like the language faculty is a system operating in accordance with its own principles, without regard for the relation between quantification (as discussed by Frege and Tarski) and truth. Expressions of natural language have the properties

they have. Speakers use those expressions, sometimes as devices for making truth-evaluable assertions. And meaning constrains use in subtle ways. But the more we learn about 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.)

5. Semantics is not Conventionalized Pragmatics

In this final section, I want isolate a point of disagreement with Stalnaker, in order to stress a more significant point of agreement. Stalnaker (1999) says that "we should separate, as best we can, questions about what language is used to do from questions about the means it provides for doing it (p.2)," and that

> The line—or more accurately a number of distinct lines—between semantics and pragmatics shift and blur. But I think there is one line that is worth continuing to draw and redraw: between an abstract account of the functions and purposes of discourse and a theory of the devices that particular languages and linguistic practices use to serve those functions and accomplish those purposes (p.16).

That seems right and important. But along with many linguists and philosophers, Stalnaker also holds that

> a principal goal of semantics is to explain how the expressions used to perform speech acts such as assertion are used to convey information—to distinguish between possibilities—and how the way complex sentences distinguish between possibilities is a function of the semantic values of their parts (2003, p.172).

I think this obscures the line worth drawing. For it conflates, under the heading 'semantics', questions about the use of natural language with questions about the nature of

linguistic expressions. It presupposes not only that a *sentence* can distinguish between possibilities, but that this sentential property is compositionally determined. And it isn't clear that these are genuine explananda—facts to be explained, as opposed to simplifying presuppositions whose falsity can be ignored for certain purposes—much less explananda for theories of the *devices* that natural language provides. There are certainly facts about what speakers can use sentences to do; and use is severely constrained by aspects of meaning that are compositional. But if semantics is primarily about what *expressions* mean, then semanticists can and should abstract from various details concerning how expressions are used. (And while certain Gricean principles may apply to discourse among rational beings, regardless of what their expressions mean, there are also principles governing what expressions of natural language cannot mean regardless of how they are used.) One should not insist that appeal to kontents explain the fact that propositions are structured in ways that sets are not. But equally, one should not insist that meaning be related to use in certain ways.

In stressing another valuable distinction, between "descriptive" and "foundational" semantic theories, Stalnaker (2003) says that the former specify Semantic Values for expressions of a language "without saying what it is *about the practice of using that language* that explains why" the expressions have those Values (p. 166, my italics). Here, he is presumably thinking about conventional facts: 'chien' signifies dogs as opposed to cats, etc. I would just omit the italicized phrase, which isn't needed to make the distinction. Though I wouldn't quibble about this, if remarks about use were not so regularly combined (in philosophy and elsewhere) with the suggestion that semantics is fundamentally contingent/arbitrary/conventional/learned. Stalnaker says that we should "all agree that it is a matter of contingent fact that the expressions of natural

language have the *character and* content that they have (p.195, my italics)." But is it obviously a contingent fact that 'John is eager to please' has the semantic character it does? This sentence is *not* apt for use as a way of saying that John is eager for us to please him. Prima facie, this negative fact is due to deep properties of the sentence, not superficial features of how the words are related to features of the environment. And perhaps many aspects of the actual character of a sentence are essential to it, much as the actual atomic number of gold is essential to it.

I'm not sure it even makes sense to think about an expression of *natural* language as having a meaning that would violate principles of Universal Grammar. (Is this just to imagine a brain that associates certain *signals* with interpretations in a nonhuman way?) And one can avoid tendentious conceptions of expressions, while endorsing what Stalnaker (2003) takes to be a central aspect of Kripke's thought—viz., that the contents of speech and thought are determined by the things with which speakers and thinkers interact.[13] We can define a useful and externalistic notion of kontent to capture this idea. But we shouldn't assume that linguistic meaning can be determined in like fashion. One is free to define 'meening' and 'kontext' so that the meening of a sentence is a function from kontexts to kontents. But then one can't assume that an adequate theory of meening, whatever that would be, is an adequate theory of understanding. Understanding, in so far as we can have theories of it, may have more to do with the language faculty than with kontents. We can define 'outerstanding' in terms of kontents. But externalist stipulation must come to an end at some point.

Stalnaker (1999) says, "The attempt to do semantics without propositional content is motivated more by pessimism about the possibility of an adequate account of propositions than it is by optimism about the possibility of explaining the phenomena without them (p.3)." This may

be true of many philosophers following Quine and Davidson. But Chomsky is optimistic about the possibility of explaining semantic phenomena as phenomena that reflect the nature and operation of the language faculty—and hence, as phenomena unlikely to be explained in terms of propositions. With this in mind, I have tried to argue here that a Chomskyan internalism about semantics is compatible with the following idea: Stalnaker's account of propositions is adequate *for purposes requiring appeal to propositions, as opposed to meanings*.

Of course, this makes no sense if we assume that "Syntax studies sentences, semantics studies propositions" (Stalnaker [1999, p.34]). Fodor, along with many others, combines this common assumption with the idea that propositions are linguistically structured. Stalnaker combines it with an opposing conception of propositions, as sets of possible worlds, and is led to say that the subject of semantics "has no essential connection with languages at all, either natural or artificial" (p.33). Fodor is led to say that expressions of natural language are themselves essentially meaningless.[14] But another possibility is to reject the slogan "semantics studies propositions," leaving it open whether this a false hypothesis or an unhelpful stipulation.

Prima facie, semantics is the study of linguistic meaning, whatever that turns out to be. If it turns out that sentences of natural language are not systematically associated with propositions, whatever *they* turn out to be, we can conclude that natural language semantics is a branch of human psychology concerned with the language faculty. We are free to invent an enterprise, Psemantics, defined as the study of how *sentences* of spoken languages (as opposed to certain human *actions*) are related to: sentences of Mentalese, kontents, Russellian propositions, or whatevers. Inquirers are free to puruse this enterprise instead of, or in addition to, studying the language faculty (independent of prior assumptions about how it is related to propositions). But if

fixation on Psemantics keeps getting us into trouble, that may be nature telling us something.

Appendix: Diagonals Obscure Nontrivial Aspects

Suppose you do take semantics to be the study of propositions, which you take to be sets of possible worlds, and you grant that 'Hesperus is Phosphorus' and 'Hesperus is Hesperus' differ in meaning. If you also suspect (*pace* Kripke) that any truth knowable *a priori* is in some sense necessary, you may be tempted to identify linguistic meanings with (functions from contexts to) "diagonal" K-propositions described in terms of the apparatus of "two dimensional" modal logic; see Chalmers (1996), Jackson (1998). Stalnaker's (2003) good arguments to the contrary, which illustrate the distinction between descriptive and foundational semantic theories (see also Kaplan [1989]), are independent of Chomsky-style considerations about negative facts. But appeal to diagonal kontents can seem like an attractive way to preserve a unified conception of semantics and pragmatics. So it is worth being clear that this only hinders discussion of natural language constraints on signal-meaning association.

Let w1, w2, and w3 be distinct possible worlds in which the sentence 'You saw us' is used to make an assertion in accordance with the following matrix.

	w1	w2	w3
w1	T	F	T
w2	F	T	F
w3	T	F	F

Suppose that at w1, Chris is talking to Pat, with Hilary as the (only) other relevant party. Then the assertion is true iff Pat saw Chris and Hilary; and it is true at w1 and w3, but not w2. Let w2 be a world at which Pat is talking to Chris, with Hilary as the other relevant party; and let w3 be a world at which Pat is talking to Chris, with the other relevant party being a fourth person, Sam.

Then at w2, Pat uses ‘You saw us’ to make an assertion that is true iff Chris saw Pat and Hilary; and at w3, Pat uses this sentence to make an assertion that is true iff Chris saw Pat and Sam. The matrix represents a function **M**, which Stalnaker (1978) calls a “propositional concept,” from worlds to sets of worlds: **M**(w1) = {w1, w3}; **M**(w2) = {w2}; **M**(w3) = {w1}. Let the “diagonal” of **M**, $D^{\mathbf{M}}$, be {w: w ∈ **M**(w)}. In our simple example, the diagonal is {w1, w2}.

Now consider the matrix below, this time taking w1, w2, and w3 to be possible worlds in which the relevant speaker uses ‘I am here’ to make an assertion.

	w1	w2	w3
w1	T	F	F
w2	F	T	F
w3	F	F	T

Suppose that at w1: Chris is at location L and speaking to Pat, who is at location L*. Then at w1, the assertion is true iff Chris is at L; and the kontent of this assertion, indicated with the first line, includes w1. Let w2 be a world where Pat is speaking to Chris from a third location L**, while Chris is at L*. Then the assertion is true iff Pat is at L**; and its kontent includes w2. Let w3 be a world where Pat is speaking to Chris from L, while Chris is at L*. Then the assertion is true iff Pat is at L; and its kontent includes w3. Call this matrix **CAP**, to highlight the contingent a priori status of the assertions. The diagonal, $D^{\mathbf{CAP}}$, is the “universal” set {w1, w2, w3}.

This provides a theoretical model of a phenomenon that might otherwise seem puzzling: when a speaker says ‘I am here’ (in typical circumstances), she *can’t* be wrong, even though she is making a claim that is only *contingently* true. Thus, $D^{\mathbf{CAP}}$ can reflects something that uses of ‘I am here’ have in common; and this is worth noting. But one shouldn’t reason as follows: the meaning of a sentence is constant across contexts; so any (semantically relevant) property of a

sentence S constant across contexts is a good candidate for being the meaning of S.

Let's recycle the first matrix, using it to evaluate 'You are easy to please'.

	w1	w2	w3	
w1	T	F	T	w1: Chris talking to Pat; other relevant party, Hilary
w2	F	T	F	w2: Pat talking to Chris; other relevant party, Hilary
w3	T	F	F	w3: Pat talking to Chris; other relevant party, Sam

The assertion at w1, true at w1 and w3, is true iff it is easy for Chris and Hilary to please Pat. The assertions at w2 and w3, made by Pat, are true (respectively) iff: it is easy for Pat and Hilary to please Chris; and it is easy for Pat and Sam to please Chris. But focus on two things that Chris *cannot* do by saying 'You are easy to please' at w1. First, Chris cannot make an an assertion that is true iff it is easy for Pat and Hilary to please Chris; although Chris could make such an assertion wth 'I am eager to please'. Second, Chris cannot make an assertion that is true iff it is easy for Pat to please Chris and Hilary; although Chris could make such an assertion with 'You can easily please (us)'. While the second fact is more theoretically interesting, appeal to diagonals at best helps to characterize the first.

Given a language like English except that the sounds of 'you' and 'I' signify first and second personal pronouns, respectively, Chris could use the sound of 'You are easy to please' to make an assertion that Chris cannot make with the homophonous English sentence. But this just illustrates a conventional (and theoretically uninteresting) feature of English. By contrast, Chris would need to speak a nonhuman language in order to use the English words as a way of saying that the addressee can easily please relevant parties. Put another way, if we represent sentential meanings with propositional matrices, we must associate 'You are easy to please' and 'It is easy for relevant parties to please you' with the *same* matrix, while associating 'It is easy for you to

please relevant parties' with a *different* matrix. This last sentence should be associated with the matrix below, given how things are at w1, w2, and w3.

	w1	w2	w3
w1	F	T	F
w2	T	F	T
w3	F	T	T

This raises the question of *why* 'You are easy to please' has the matrix it does, as opposed to the one immediately above. And the force of this question is heightened when we note that the superficially similar sentence 'You are eager to please' has a different semantic character, according to which any an assertion made by using it is true iff the addressee is eager to be one who pleases relevant parties. (Indeed, it is easy to describe the worlds so that 'You are eager to please' corresponds to the matrix immediately above; though at this point, few readers will be eager for me to do so. Note also that 'The goose is ready to eat' is ambiguous.)

The idea of identifying linguistic meanings with diagonals seems hopeless once one considers the semantic relations among sentences like 'John persuaded Bill to leave', 'Bill intended to leave', and 'John expected Bill to leave'. With enough effort, one can probably specify an algorithm according to which: the second is true at every world where the first is true, while false at some but not all worlds where the second is true; and each sentence can used to make all and only the assertions it can be used to make. But why think that such an algorithm would be characterizing the meanings of the sentences, as opposed to describing certain features of the assertions one could make with independently meaningful sentences? At best, appeal to diagonals provides a way of capturing certain trivial examples of a priori "knowledge" corresponding to contingent/conventional aspects of natural language semantics. The real work

lies with describing more interesting aspects of meaning in theoretically perspicuous ways; see Chomsky (1965).[15]

References

Baker, M. (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).

Boolos, G. (1998). *Logic, Logic, and Logic* (Cambridge, Mass: Harvard University Press).

Cappellen, H. and Lepore, E. (forthcoming). *Insensitive Semantics* (Cambridge, Mass: Blackwell).

Carnap, R. (1950). 'Semantics, Empiricism, and Ontology', *Revue Internationale de Philosophie* 4: 208-28.

Carruthers, P. (2002). 'The Cognitive Functions of Language', *Behavioral and Brain Sciences* 25:261-316.

Chalmers, D. (1996). *The Conscious Mind* (Oxford: Oxford University Press).

Chomsky, N., (1957). *Syntactic Structures* (The Hague: Mouton).

---------(1965). *Aspects of the Theory of Syntax* (Cambridge, Mass: MIT Press).

---------(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.

---------(1995). *The Minimalist Program* (Cambridge, Mass: MIT 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, Mass: MIT Press).

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.

Crain, S., Gualmini, A. and Pietroski, P. (forthcoming). ''

Davidson, D. (1967a). 'Truth and Meaning', *Synthese* 17: 304-23.

---------(1967b). 'The Logical Form of Action Sentences', in N. Rescher, N. (ed.)
The Logic of Decision and Action (Pittsburgh: University of Pittsburgh Press).

---------(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).

Dummett, M. (1975). What is a Theory of Meaning? In S. Guttenplan (ed.), *Mind and Language*.
(Oxford: Oxford University Press).

Evans, G. and McDowell, J. (eds.) (1976). *Truth and Meaning*
(Oxford: Oxford University Press).

Fodor, J. (1975). *The Language of Thought* (New York: Crowell).

---------(1978). 'Propositional Attitudes', *The Monist* 61:501-23.

---------(1987). *Psychosemantics* (Cambridge, Mass: MIT Press).

---------(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).

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. (1985). 'On Semantics', *Linguistic Inquiry* 16: 547-93.

---------(1986). 'Lingustic Theory and Davidson's Program', in E. Lepore (ed.), *Truth and Interpretation* (Oxford: Blackwell).

---------(1989). 'Elucidations of Meaning', *Linguistics and Philosophy* 12: 465-517.

---------(1994). 'Priorities in the Philosophy of Thought', *Proc. Aristotelian Society* supp. vol, 20: 85-106.

Hornstein, N. (1984). *Logic as Grammar* (Cambridge, Mass: MIT Press).

Hornstein, N. and Lightfoot, D. (1981). *Explanation in Lnguistics* (London: Longman).

Horty, J. (ms). *Frege on definitions: a case study of semantic content* (University of Maryland).

Horwich, P. (1997). 'The Composition of Meanings', *Philosophical Review* 106: 503-32.

Jackendoff, R. (1990). *Semantic Structures* (Cambridge, Mass: MIT Press).

---------(1993). *Patterns in the Mind*.

Jackson, F. (1998). *From Metaphysics to Ethics: A Defense of Conceptual Analysis*

(Oxford: Oxford University 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.

Kripke, S. (1980). *Naming and Necessity* (Cambridge, Mass: Harvard University Press)

Larson and Segal (1995). *Knowledge of Meaning* (Cambridge, Mass: MIT Press).

Laurence, S. and Margolis, E. (2001). 'The Poverty of Stimulus Argument', *British Journal for the Philosophy of Science* 52: 217-276.

Lewis, D. (1972). 'General Semantics', in D. Davidson and G. Harman (eds.), *Semantics of Natural Language* (Dordrecht: Reidel).

---------(1986). *The Plurality of Worlds* (Oxford: Basil Blackwell).

Ludlow, P. (2002). 'LF and Natural Logic', in Preyer and Peters (2002).

Matthews, R. (2003). 'Does Linguistic Competence Require Knowledge of Language?' in Barber (2003).

McDowell, J. (1976). 'Truth Conditions, Bivalence and Verificationism', in Evans and McDowell (1976).

McGilvray, J. (1999). *Chomsky: Language, Mind and Politics* (Cambridge: Polity Press).

Montague, R. (1974). *Formal Philosophy* (New Haven: Yale University Press).

Moravscik, J. (1975). *Understanding language* (The Hague: Mouton).

Neale, S. (1990). *Descriptions* (Cambridge, Mass: MIT Press).

---------(1993). 'Grammatical Form, Logical Form, and Incomplete Symbols' in A. Irvine & G. Wedeking (eds.), *Russell and Analytic Philosophy* (Toronto: University of Toronto).

Parsons, T. (1990). *Events in the Semantics of English* (Cambridge, Mass: MIT Press).

Peacocke, C. (1999). *Being Known*. (New York: Oxford University Press).

Pietroski, P. (1998). 'Actions, Adjuncts, and Agency', *Mind* 107: 73-111.

---------(2000). *Causing Actions* (Oxford: Oxford University Press).

---------(2003a). 'The Character of Natural Language Semantics' in Barber (2003).

---------(2003b). 'Quantification and Second-Order Monadictity',
Philosophical Perspectives 17:259-98.

---------(2005a). *Events and Semantic Architecture* (Oxford: Oxford University Press).

---------(2005b) 'Meaning Before Truth'. In Preyer and G. Peters (2005).

Preyer, G. and G. Peters, G. (eds.) 2002: *Logical Form and Language*
(Oxford: Oxford University Press).

---------(2005). Contextualism in Philosophy (Oxford: Oxford University Press).

Quine, W.V.O. (1960). *Word and Object* (Cambridge Mass: MIT Press).

Ross, J. (1967). *Constraints on Variables in Syntax* (Doctoral dissertation, MIT).

Schein, B. 1993: *Events and Plurals* (Cambridge, Mass: MIT Press).

---------(2002). 'Events and the Semantic Content of Thematic Relations', in
Preyer and Peters (2002).

---------(forthcoming). *Conjunction Reduction Redux* (Cambridge, Mass: MIT Press).

Soames, S. (2002). *Beyond Rigidity* (Oxford: Oxford University Press).

Spelke, E. (2002). 'Developing knowledge of space: Core systems and newcombinations' in
S. Kosslyn and A. Galaburda (eds.), *Languages of the Brain*
(Cambridge, MA: Harvard University Press).

Stalnaker, R. (1970). ‘Pragmatics’, *Synthese* 22: xx-yy. Reprinted in Stalnaker (1999: 31-46).

---------(1976). ‘Possible Worlds’ *Nous* 10: xx-yy.

---------(1978). ‘Assertion’, *Syntax and Semantics* 9: 315-32.

Reprinted in Stalnaker (1999: 78-95).

--------(1984). *Inquiry* (Cambridge, Mass: MIT Press).

---------(1999). *Context and Content* (Oxford: Oxford University Press).

---------(2003). *Ways a World Might Be* (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.

Strawson, P. (1950). ‘On Referring’, *Mind* 59: 320-44.

Tarski, A. (1983). *Logic, Semantics, Metamathematics* (Indianapolis: Hackett).

Tenny, C. (1994): *Aspectual Roles and and the Syntax-Semantics Interface* (Dordrecht: Kluwer).

Travis. L. (1984): *Parameters and the Effects of Word-Order Variation*

(Doctoral Dissertation, MIT)

Notes

1. See Kripke (1980, pp. 15-20), who also says that his third lecture 'suggests that a good deal of what contemporary philosophy regards as mere physical necessity is actually necessary *tout court*'. For discussion in the context of supervenience theses, see Pietroski (2000, chapter six). Stalnaker's (2003) notion may be a little broader, since he speaks of ways *a* world might be (though see also p. 215); and he says the *concept* of possibility is to be understood functionally, "as what one is distinguishing between when one says how things are" (p.8).

2. Theorists employing such a notion would have to say what π-Worlds are, in enough detail to support their proffered explanations. But perhaps appeal to Ludwigian totalities, or "ersatz" analogs that are formally similar in certain respects, will be useful here.

3. Or a little more precisely (using '#' as a corner-quote), one need not endorse the following generalization: #it is logically possible that P# is true iff the semantic correlate of #that P# is a nonempty K-proposition. See Peacocke (1999) for related discussion.

4. In my view, gestures towards (remotely plausible) causal/functional-role/teleological theories are just that. Nobody has a good *theory* of why 'chien' stands for dogs as opposed to cats.

5. Compare Dummett (1975), McDowell (1976).

6. If only for simplicity, let's assume that interpretations are constant across languages, and let's pretend that all speakers of a "named language" like English have the same I-language.

7. While 'Begriffsplan' is compact, 'Begriffskonstruktionsanleitung' (concept-construction-instruction) displays the point even better. Within analytic philosophy, it was long held that expressions of natural language can't be systematically associated with (A-propositions) or concepts. But given developments in the study of language and logic, it now seems clear that the

most famous cases alleged mismatches between grammatical form and logical form were misdiagnosed, that the resources for rediganosing other cases are considerable, and that positing such mismatches is problematic (especially in light of how children can understand complex constructions). For review, see Pietroski (2003b); see also Neale (1990, 1993), Ludlow (2002).

8. I endorse the former but not the latter. In my view, Katz and Fodor (1963) rightly eschewed Lewis's (1972) stipulations about what semantics must be. But my claim is not that meanings are concepts. Recall that Strawson (1950) urged us to characterize the meaning of a referential device R in terms of "general directions" for using R, *and not* in terms of some entity allegedly denoted by R.

9. See note 7; cf. Jackendoff (1990). One does not need the mismatch hypothesis to say that in many contexts, a speaker who uses sentence Σ to make an assertion is also "entertaining" one or more sentences $S_1...S_n$, and that the truth or falsity of the speakers utterance (of Σ) depends—in complicated ways that frustrate attempts to provide theories of truth for natural languages—on the meanings of Σ *and* $S_1...S_n$.

10. See for example, Hornstein and Lightfoot (1981), Jackendoff (1993), Baker (2001), Crain and Pietroski (2001, 2002), Laurence and Margolis (2001), Crain et.al. (forthcoming).

11. Though in my view, Davidson was importantly right about the basic *structure* of semantic theories, the need for "event" variables, and the use of an extensional metalanguage; see Higginbotham (1985, 1986), Larson and Segal (1995).

12. See Pietroski (2003a, forthcoming), drawing on Chomsky (1977, 2000a), for discussion of examples like the following: France is hexagonal, and it is a republic; the red book is too heavy, though it was favorably reviewed, and the blue one is boring, though everyone is reading it; 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. See also Moravscik (1975), Hornstein (1984), McGilvray (1999).

13. One can also agree that meanings, as abstracta, are not in the head; see Stalnaker (2003, pp. 204-210).

14. Soames (2002) is an interesting case deserving separate treatment. But given his views about propositions, Soames is led to say (in the absence of confirming evidence) that *sentences* carry information, that 'Hesperus is bright' and 'Phosphorus is bright' carry the *same* information, and that these sentences are synonymous. Though Soames may be right that this is the best option, all things considered, for those who take semantics to be the study of propositions and how they are related to sentences.

15. This paper has roots in a series of conversations with my thesis supervisor, to whom I owe much. The intervening years have only deepened my appreciation of Bob's work as a philosopher and teacher. My thanks also to Dan Blair, Susan Dwyer, and Norbert Hornstein for helpful comments on earlier versions of this material.