

KENT JOHNSON*

TACIT BELIEF, SEMANTICS AND GRAMMAR

ABSTRACT. This paper explores speakers' epistemic access to the semantic and syntactic features of sentences of their language. I argue that there is evidence that *ceteris paribus*, the actual semantic features of sentences of a language are accessible as such by typical speakers of that language. I then explore various linguistic, cognitive, and epistemic consequences of this position.

0. INTRODUCTION

It is standard in both linguistics and philosophy to assume that we have tacit belief in a grammar of our language.¹ By assuming that we bear a cognitive relation like belief (or “cognizing”, cf. Chomsky 1986, p. 265) to a grammar, we can explain the cognitive aspects of our linguistic abilities. But by supposing that this belief is only tacit held, we can explain why typical speakers are not consciously aware of so many of the linguistic features that serve to structure their language. Even in relatively simple circumstances, the usefulness of the notion of tacit belief is readily apparent. Consider, for example, sentence (1):

- (1) The rabbit gave the elephant the lion.

Typical speakers have no difficulty understanding (1). They know (and thus truly believe), for instance, that it expresses that the elephant received the lion, and not the other way around. On the other hand, speakers are much less aware of how (1) is structured. Ask a typical speaker of English how sentence (1) is put together, and you won't receive much more detail than that *the rabbit* is the subject, *gave* is the verb with *the elephant* and *the*

* I have received very useful criticism and feedback from many people, including Mark Baker, Jim Higginbotham, Ernie Lepore, Robert Matthews, Robert May, Mark Moyer, Terry Parsons, Paul Pietroski, Karin Stromswold, Jonathan Sutton, Wayne Wright, and two anonymous reviewers.

¹ This proposal has been discussed by e.g., Chomsky (1986, pp. 263-273); cf. also Chomsky (1965), Fodor (1963), Graves et al. (1973), Davies (1981), Evans (1981), Pinker (1979, 1993), Lycan (1986).



lion as its direct and indirect objects. However, current linguistic theory strongly supports the view that there is an entire microcosm of syntactic activity in the verb phrase of (1). Indeed, a popular view in linguistics is that there are actually two verbal structures present in (1), although only one of them (*gave*) is pronounced, and that the configurational properties of the unpronounced verb is much like the structure of a passive verb with *the elephant* as its subject (e.g., Larson 1988). Assuming that such an account is true, it's a little surprising that speakers are so unaware of the structure of the sentence, especially since they have a strong grasp of the sentence's meaning. After all, the sentence's structure is part of what gives meaning to the sentence. In this paper, I will explore the issue of whether the distinction between tacit and non-tacit awareness of linguistic features is correlated with the distinction between syntactic and semantic features of language. The philosophical literature has given some attention to this issue. Many philosophers tend to think that we do always grasp the meaning of at least our own utterances (e.g., Peacocke 1998, p. 68; Burge 1990, p. 31, fn. 4; Schiffer 1987, pp. 255–267; cf. Dummett 1974). In fact, there is a tradition going back to Wittgenstein which suggests that it is both necessary and *a priori* that we know what our own utterances mean (cf. Wright 1986, p. 219). But there are other philosophers who have been inspired by the complexity of the linguistic data, and they have held that we do not always grasp all the semantic properties of our utterances (e.g., Higginbotham 1998, p. 152, 1989, p. 159; Larson and Segal 1995, pp. 542–553). This latter view is also commonly held by linguists. Although I am sympathetic to the methodology that supports this latter view, in this paper I will argue that the former view may in fact be correct in some important respects. The argument suggests that the semantic and syntactic components of language differ epistemologically in that only the latter are tacit, but the former are graspable by our reasoning minds. However, we will also see that this view is limited in scope, and that the true epistemological situation is somewhere in the middle.

The thesis that (roughly speaking), syntax is tacit but semantics is not is itself a substantial claim about the epistemological organization of language. We will also see that features of this claim interact with various views present in the literature regarding the modular nature of language and the construction of linguistic theories. The thesis also supports a number of philosophically important claims, two of which I will discuss in this paper. In the first place, there has been a substantial amount of discussion about whether we think in English. Intuitively, this view seems plausible, especially for certain simple cases, like when you entertain a belief such as “this morning for breakfast, I had . . .” Thinking in these cases seems

to be a form of “speaking silently to oneself”. Despite the plausibility of this view, however, I will argue that my main thesis provides reason to doubt that we think in our natural language. Secondly, the thesis helps us to understand the public nature of language. Some of the most central debates in philosophy of language and philosophy of mind concern the question of our epistemic relation to our public language and our apparent ability to understand the speech of another member of our linguistic community. I will argue that the thesis explains how we understand the general structure of the atomic formulas of our public language.

Although I am primarily interested in issues of contemporary concern, the methodology I will use is older. It has its origins in the work of J. L. Austin, particularly in “A Plea for Excuses”. There Austin articulated a method for studying how we categorize things and events in the world by studying how we use language to describe them. Using a variety of examples, he showed that there are numerous subtle conditions relevant to whether two words of the same general sort (verb, adverb, noun, etc.) can both appear in the same places. His examples showed that often they cannot. Austin was keenly aware of the importance of these differing distributions. He writes that “we are sometimes not so good at observing what we *can't* say as what we can, yet the first is pretty regularly the more revealing” (Austin 1956, p. 16).² Contemporary theoretical linguistics has confirmed Austin’s observation: we can learn a lot about the meanings of words and classes of words by examining their differing patterns of grammatical distribution. The key to finding evidence that syntax and semantics differ epistemologically involves studying the details of some patterns of distribution of some *prima facie* similar words. These distributional patterns reveal a rich semantic structure buried underneath a superficially transparent structure of simple sentences. I will suggest that this underlying semantic structure is accessible by us. The accessibility of this relatively “hidden” semantic structure is surprising, because syntactic structures which are fairly “superficial” in the language appear to be only tacitly grasped.

² For instance, he uncovers numerous subtle distributional patterns regarding adverbs that might be used to modify a verb phrase so as to form an excuse. An example of this appears in his discussion of *inadvertently*:

In passing the butter I do not knock over the cream-jug, though I do (inadvertently) knock over the teacup – yet I do not by-pass the cream jug *advertently*: for at this level, below supervision in detail, *anything* that we do is, if you like, inadvertent, though we only call it so, and indeed only call it something we have done, if there is something untoward about it. (Austin 1956, p. 18)

The paper is structured as follows. In section 1, I discuss some salient properties of tacit belief and develop a taxonomy of cognitive states around it. In section 2, I take a simple syntactic feature and construct a hypothetical test for its epistemic status. We will see there that even very simple syntactic features are tightly restricted in terms of the ways that typical speakers can employ them. In section 3, I describe some relatively obscure semantic features of language and argue that we are justified in supposing that they really exist. In section 4, I construct a hypothetical test for the epistemic status of one of these semantic features. I will then argue that the two hypothetical tests support the thesis that there is a marked epistemic distinction between syntax and semantics. I develop the theory and respond to some objections in section 5. In section 6 I argue that the thesis contributes substantially to two philosophical issues. I conclude in section 7.

Before beginning, three caveats are in order. First, I will not question whether we have some sort of knowledge or belief in a grammar. Some writers have suggested that our “knowledge of language” is actually only a complex disposition to transmit (or “transduce”) information between a language of thought and our speech perception and production mechanisms (e.g., Fodor 1975, 1990, 1998; Schiffer 1987; Evans 1981; Lepore 1996; for related issues of another sort, cf. Pettit 2002). All of my main points in this paper will hold *mutatis mutandis* for such theories. At heart, much of the present project is epistemological, and matters of broad cognitive structure will not matter much. (Higginbotham has argued in several places that the sort of view advocated by Fodor and Schiffer has little effect on linguistics and the philosophy of linguistics, including the issues to be discussed here; cf. Higginbotham 1988, 1994, 1995, 1998.) With that said, I will continue to assume that we have knowledge of a grammar of our language. The second caveat is that I will be assuming a relatively general form of GB syntax in what follows (Chomsky 1981). I choose this type of syntactic theory because it is by far the most common. However, most of the discussion will be concerned with patterns of data rather than the structure of any particular theory, so I believe (although I will not argue for it) that the claims I make can be translated into similar claims for other versions of syntax.³ The final caveat is that language here should be understood internalistically (cf. Chomsky 2000, ch. 7): a language of a person is that state of a speaker’s mind that enables them to acquire and speak the language they do. Towards the end of this paper, I will address

³ For recent neurological evidence in favor of the particular syntactic assumptions I make, cf. Feathersen et al. (2000).

some issues about language as a public phenomenon, but I will usually be thinking of language as a kind of cognitive capacity.

1. TACIT BELIEF

Since we are interested in the boundary between the tacit and the non-tacit, we need to begin by clarifying what is to count as tacit belief. Roughly speaking, in philosophy and elsewhere the most interesting distinction is between those beliefs that are in some sense “available” or “accessible” for use in reasoning, decision-making, and reflection, and those beliefs that are not (e.g., Stich 1978; Lycan 1986; Davies 1989; Kirsh 1990; Johnson and Lepore 2003). This distinction will be a primary tool for classifying beliefs, so it will be useful to introduce some labels now. I distinguish three kinds of belief. On the one hand, there are *occurrent* (or explicit) beliefs, which are what you have when you consciously entertain a proposition you believe. If you are asked to calculate two plus two, your conscious belief that two plus two is four is occurrently believed. You spend much of your life not occurrently believing that two plus two is four, but that belief is readily available to you, should the need arise. If a belief is occurrent or at least available (in the relevant sense) to a person, I will call it *accessible* by that person. Accessible beliefs are those that we can use in reasoning, decision-making, and reflection, in some interesting sense of *can*. Of course, not anything goes here: if you have to look something up in a book to form the occurrent belief that P, then you don’t have an accessible belief that P. *Tacit* beliefs, on the other hand, are those beliefs that cannot be used in these activities. Using these labels, the current project is to distinguish between the features of language that are accessible from those that are only tacitly grasped. I will assume that a linguistic feature is accessible iff it is part of the content of accessible beliefs, in a sense to be filled out below; tacit linguistic features are those features that only figure into tacit beliefs.⁴

A glance at the relevant literature displays three important properties regarding accessible beliefs. First, the distinction between tacit and accessible beliefs is vague.⁵ It is not accidental that many theories about tacit belief leave a large gray area in between the definitely tacit and the definitely accessible. I suspect that the vagueness in the distinction is inherent,

⁴ Not every discussion of tacit knowledge and belief centers around this issue. For instance, Crimmins (1992) and Lycan (1986) both focus more on the difference between occurrent and non-occurrent beliefs.

⁵ E.g., Stich (1978, p. 500), Higginbotham (1998, pp. 151–152), Dennett (1978, pp. 46–48); but cf. Higginbotham (1998b, pp. 439–440), Lycan (1986), Kirsh (1990).

so it is unsurprising that a precise distinction between the tacit and the accessible is hard to find. Second, even though accessible beliefs can be used in explicit (i.e., occurrent) reasoning, explicitly entertaining them can produce unexpected results for the person. For instance, many speakers will misinterpret (2):

- (2) No eye injury is too trivial to ignore.

It is easy to understand the sentence as meaning that no eye injury is so trivial that one should ignore it. However, a bit of reflective computation shows that (2) actually means that no eye injury is so trivial that one cannot ignore it (Higginbotham 1988, p. 237, fn. 19). Does this mean that our belief in the correct meaning of (2) is only tacit? It does not appear to be, since a moment's thought yields the correct meaning; thus there is a good sense in which the belief is available for use by us. The third feature of accessible beliefs is that to become occurrent, they may require some kind of priming influence. If you misinterpreted (2), you might not have explicitly arrived at a correct interpretation of it unless you were prompted to do so by being told what the sentence doesn't mean. Similarly, unless you are asked whether the sentence *Do geese see God?* is a palindrome, you might not have an occurrent belief that it is, but once you are asked, it is easy to see that it is. What counts as legitimate priming is an important issue. In fact, a theory of legitimate priming is probably tantamount to a theory of accessible belief. To see this, consider a case of illegitimate priming. Suppose a logic teacher tells a struggling student that Ramsey's theorem follows from the compactness theorem, and the student comes to believe this, even though she had never thought about it before. Clearly this is a case where the "priming" that generates the occurrent belief in the student provides no evidence at all that the student had an accessible belief that Ramsey's theorem follows from compactness. Similarly, asking a student whether she believes that the answer to the problem is 16 may not be an instance of priming but of hinting or just plain teaching (cf. Stich 1978, pp. 505–506 for an interesting example from cognitive psychology). In short, a theory of priming – and thus a theory of accessible belief – must address a question as old as Plato's *Meno*: What counts as learning something new, and what counts as simply being brought to attend to something you already know?

Is there anything we can say about what counts as legitimate priming? Crucially, we want to distinguish circumstances where the person learns something new from circumstances where she is primed to attend to something she already knows. There seem to be at least two marks that suggest that one already believes a given proposition. The first is

the readiness with which one comes to occurrently grasp and entertain the proposition in question. If only a little priming enables the subject to correctly feel that she understands what is being talked about, then this is evidence that she has an accessible grasp of the proposition in question. A second mark of accessibility concerns the subject's ability to correctly draw novel inferences involving the concept. If you are only told a little bit about some claim, and you are then able to use it in tasks that are distinct from the circumstances used to describe it, then that is evidence that you accessibly grasp the proposition. (Peacocke (1998) has recently explored this phenomenon from a purely philosophical perspective; cf. also Higginbotham 1998b, pp. 151–152.) Despite the unclarity of so many aspects of the tacit/accessible distinction, these two features of legitimate priming will be useful in helping us to distinguish those features of language that are accessible from those that are merely tacit. I turn now to the case of syntax.

2. SYNTACTIC FEATURES AND ACCESSIBILITY

The goal of this section is to take a simple syntactic feature and construct a hypothetical test to see if it is accessible or tacit. The feature I will explore is called a “(wh-)trace”. Traces occur in questions, where a *wh*-word is moved from its normal position to the front of a sentence. Thus, we see the presence of traces in the questions in (3).

- (3) a. Melanie kicked John.
a'. Who did Melanie kick *t*?
b. You went to Boston.
b'. Where did you go *t*?

In (3a), we have a simple transitive verb. In (3a'), the object of the verb has been replaced with *who*, and for various linguistic reasons, the *wh*-word moves to the front of the sentence. Similarly, the prepositional phrase in (3b) is replaced with *where*, which also moves to the front of a sentence. Traces are denoted by *t* in (3a') and (3b'); they serve to mark the position in the sentence that the *wh*-word is associated with. We can get some idea of why syntacticians might have wanted to posit an unpronounced

feature in the syntax of the sentence if we look at some simple patterns of ungrammaticality, like those given in (4).

- (4) a. *Melanie kicked John Bill.
 a'. *Who did Melanie kick Bill?
 b. *You went to Boston to Cairo.
 b'. *Where did you go to Cairo?

(4a) is ungrammatical because the verb *kick* has too many objects. But similarly, (4a') is also ungrammatical, and it is not implausible to suppose that this is because *kick* still has too many direct objects. Similarly, (4b) is ungrammatical, because there are too many specifications of the location towards which the going was directed.⁶ I do not pretend to have given a complete justification of the existence of traces in English syntax. In fact, part of my present claim is that it is very hard, short of actually studying syntax, to grasp many features that are nevertheless plausibly present in the structure of sentences. (Notice, incidentally, that I characterized traces semantically, in terms of their effects on the logical structure of clauses. I did this because it appears to be the easiest way to introduce such features; I know of no similar way to quickly introduce them in syntactic terms alone.)

I will take the existence of traces in many sentences for granted.⁷ I want to explore whether we have a tacit or accessible grasp of this linguistic feature. As we saw in section 1, the accessible features are the ones that “can” be used in occurrent reasoning even though typical speakers may need some priming with respect to the feature in order to determine whether they can use it or not. So I will describe an experiment that would determine whether speakers can use the feature after they have been suitably primed. Speakers could be primed about traces by being given sentences like those in (3), and the notion of a trace could be justified by means of data like that in (4). Obviously, though, legitimate priming precludes a complete course in syntax: restricting the priming they receive will allow them to display a readiness to grasp the feature in question, if it is in fact accessible. After some examples, speakers should start to feel relatively comfortable with traces. Eventually they will be able to spot the traces in simple examples

⁶ There is a tremendous amount more that could be said about this latter example. The dividing up of thematic roles in a clause lies at the heart of lexical semantics and the interaction between semantics and syntax. (For some discussion, cf. e.g., Tenny 1994; Grimshaw 1990; Jackendoff 1990.)

⁷ These assumptions are ultimately quite weak. Traces need not be constituents of sentences in the same way that nouns are. Instead, they may only be configurational relations between features. Such a notion of a trace will suffice for my purposes.

that are within the “training space” of constructions used to prime them, like those in (5).

- (5) a. What did you eat?
 b. How did you fix the car?

But being able to identify traces as they occur in examples that are so similar to the ones used to prime the speaker is hardly a display of the grasp of the feature. Our discussion of the tacit/accessible distinction suggests that if a feature is really accessible, then with only a little priming, the subject should be able to use it in novel situations. So in the present linguistic case, once speakers have received the priming data, we should test whether they can identify instances of traces as they occur in novel environments (i.e., outside of the training space). Some examples of constructions we might give them (assuming they were not used as part of the priming data) are given below in (6). We might give them some of these constructions and ask them to identify the sentences that contain traces, paying special attention to the underlined expressions.

- (6) a. John is easy to please;
 b. John is eager to please;
 c. John seems to be happy;
 d. John wants to kiss Shannon;
 e. Who do you think John kissed?
 f. Which report did you file without reading?

Assuming the data in (6) is novel to speakers, it seems clear that they will not be able to distinguish the traces in any of these constructions. Moreover, even if a speaker happened to guess that a trace occurred in a sentence like (6a), for instance, she would not be able to distinguish that trace from the distinct empty category in (6b). However, syntacticians often posit a number of unpronounced features in the various sentences, some of which are given in (7).

- (7) a. John is easy to please *t*;
 b. John is eager PRO to please;
 c. John seems *t* happy;
 d. John wants PRO to kiss Shannon;
 e. Who do you think *t* John kissed *t*?
 f. Which report *t* did you file *t* without reading *e*?

Without getting into too much detail, we can see some basic plausibility in supposing that *John* started out as the subject of the lower clause *to please*. After all, *John* plays no real semantic role as the subject of *is easy to please*, but is there simply because English (unlike other languages) demands that subjects be overtly realized. (8a) shows that we can satisfy English's desire for a subject without using *John*. However, (7b) is different: *John* is a genuine participant in the event described by the verb phrase *is eager to please*. As (8b) shows, we cannot move *John* to the lower clause.

- (8) a. To please John is easy;
 a'. It is easy to please John;
 b. *To please John is eager;
 b'. *It is eager to please John.

Ultimately, the reason why the subject of *to please* in (7b) is not just a trace of where *John* used to be is that languages don't like the same syntactic expression filling more than one semantic role. (In syntax, these semantic roles are called "thematic roles", and the restriction that each role can be filled by just one logical argument is part of the "theta criterion".) Since *John* is playing a role as subject of *eager*, that same syntactic feature cannot also play a role as the object of *to please*. (Again, these examples are only intended to be illustrative of the sorts of data that syntacticians account for with various unpronounced features. For more detailed treatments of these subjects, a syntax textbook should be consulted (e.g., Haegeman 1994; Culicover 1997).)

At this point, we seem to have a plausible explanation of the inaccessibility of traces. One possibility is that the inaccessibility of traces is due to their "theoretical" status. That is, traces may be inaccessible because their existence is justified *via* research into distributional patterns of sentences (within English and across other languages). If this rationale is correct, we should expect that if we justify hypothesizing the presence of a feature in language primarily by the theoretical work that it does, then the feature will be only tacitly grasped. While I think this explanation and the prediction it makes are both plausible, I also think that they are false. My explanation why they are false requires some of the apparatus that I will develop in the next two sections. I will return to this issue in section 5.

It's worth observing that the methodology employed in this section is a real case of letting "the abnormal throw light on the normal" (Austin 1956, p. 6). The patterns of ungrammatical sentences are at least as important as the patterns of grammatical ones. For without the ungrammatical sentences, we would have no reason to believe that such features as traces

really are present in the language, and without reason to believe traces are present, we lose reason to believe speakers are failing to see something that is there. In this sense, we are following Austin in that we are devising an epistemological “‘laboratory technique’ which could be fruitfully used for finding solutions to [sc. certain philosophical problems] very much fuller, more systematic, and more accurate than any hitherto” (Urmson 1965, p. 232). We will use this technique again in section 4, when we explore the epistemological properties of semantic features.

3. SUBLEXICAL CONCEPTS

In the last section I argued that relatively basic syntactic features like traces are not accessible. Does the same hold for semantic features? Obviously not in the simplest cases: speakers can easily tell that whether there is anything about dogs in a sentence like *the dog bit the boy*, or in one like *the ring is made of gold*. However, the question of the accessibility of semantic features becomes more interesting when we turn to more obscure ones. In this section, I will unearth some of these features. I will argue that we have good reason for taking seriously the idea that the semantic structure of a sentence often contains numerous aspects of meaning that are not associated with a whole word or with one word only. Since these features often constitute only aspects of a word’s meaning, I will call them *sublexical concepts*. After I defend the existence of sublexical concepts in this section, I will in the next section explore their epistemic status.

I begin with the well-studied phenomenon of locative verbs (cf. Levin 1993 for many references). Locative verbs are probably best introduced via an example. (9) shows that some verbs fit into both syntactic patterns listed in (10).

- (9) a. I sprayed the flowers with water.
 b. I sprayed water onto the flowers.
 c. Similar verbs: *cram, plant, smear, splash, encrust, strew, wrap, empty, load*.

- (10) a. Subject Verb Object₁ *with* Object₂;
 b. Subject Verb Object₂ *into/onto* Object₁.

But many other verbs appear in only one of these patterns. (11) shows that some verbs only fit into the form in (10a), and (12) shows that other verbs only appear in the form in (10b).

- (11) a. John drenched the flowers with water.
 b. *John drenched water onto the flowers.
 c. Similar verbs: *cover, decorate, fill, soil, bathe, saturate, infect, taint, pollute, litter, ornament, pave.*
- (12) a. *Alex threw the flowers with the water.
 b. Alex threw the water onto the flowers.
 c. Similar verbs: *pour, arrange, immerse, lodge, mount, pour, spill, coil, shove, funnel.*

(Levin 1993, pp. 50–51, and Pinker 1989, pp. 126–130 contain many more verbs in all three classifications.) What could account for this distribution of verbs? Still following Austin’s lead, we want to take these verbs and “prise them off the world, to hold them apart from and against it” (Austin 1956, p. 8). By doing this, and by “observing what we can’t say” as much as “what we can” (Ibid.), we can see that the three classes of verbs given in (9), (11) and (12) have a certain semantic similarity. Intuitively speaking, in each class, the verbs have a core meaning common to the other verbs in that class, but distinct from the core meaning of the verbs in the other classes. We can characterize these meanings roughly as follows:

- (13) a. *Alternating Verbs*: Subject changes Place by causing Stuff to go into/onto Place in a certain way.
 b. *With-Only Verbs*: Subject changes Place.
 c. *Into-Only Verbs*: Subject causes Stuff to go into/onto Place in a certain way.

To spray the flowers with water, it is not enough that one simply wets one petal of one of the flowers; one must cover all or most (or enough of) the flowers in a certain way. Similarly, the meaning of *spray* also restricts how the water goes onto the flowers: one hasn’t sprayed the flowers with water if one has only poured water onto them, or dunked the flowers in a pail full of water. In the case of verbs like *drench*, though, there is only a requirement that the flowers end up soaked with water, whether by spraying them or immersing them, or whatever. In contrast, verbs like *throw* are just the opposite. Throwing water on the flowers specifies how the water gets onto the flowers; e.g., dunking the flowers in a bucket of water would not count.

However, *throw* does not specify the end state of the flowers after the water has been thrown onto them. It would, for instance, be perfectly natural to report that after Mary threw the bucket of leftover water onto the flowers, most of the flowers were still dry.

The semantic explanations of these verb classes is robust. As new verbs in these semantic classes enter the language, they obey the grammatical patterns predicted by them. For instance, you might report that you downloaded a file onto your hard drive, but you wouldn't say that you *downloaded your hard drive with a file, since the verb *download* does not require that the site of downloading be changed (or "completely affected"; cf. Tenny 1994) by the downloading. Similarly, you might *ftp a file into your email account*, but you would not **ftp your email account with a file*.

The ability of sublexical concepts to explain the grammatical behavior of words suggests that aspects of meaning – such as whether a verb expresses that it changes a Place or that some Stuff moved in a certain way – are genuinely semantic components of language.⁸ However, as we will soon see, such sublexical concepts as CAUSE and CHANGE as they occur in (13) are not our ordinary concepts of causation and change. So when we appeal to these sublexical concepts in our explanations, we are positing components of meaning that do not correspond to the meanings of any individual words of English. At first glance, this might make it seem that the price of positing sublexical concepts is too high. Although these semantic features capture the distributions of locative verbs nicely, they are a pretty heavy-duty piece of apparatus. If all they do is account for locative verbs, then perhaps the price of positing them is not worth the work they do. This worry can be alleviated by observing that sublexical concepts do much more than account for locative verb distributions. I will briefly list four other areas where they are theoretically useful. First, sublexical concepts are frequently used to explain various other grammatical distributions and semantic properties. For instance, the notion of causation mentioned above appears to be part of accomplishment clauses. *Denise built a house* means that Denise was building a house, and a "direct" result of her performing this activity was that a house came into existence (e.g., Vendler 1967; Dowty 1979). In the next section, I will also discuss some further phenomena involving causation. Second, sublexical concepts have been applied in various areas of psycholinguistics. For instance, Gropen et al. constructed scenarios where young children could use locative verbs (cf., Pinker 1989; Gropen et al. 1991; Pinker 1999). They presented evidence that when

⁸ I will follow a common trend in using the concepts occurring in (13) (e.g., Pinker 1989; Hale and Keyser 1987), although the general argument I will construct using them has fairly easy correlates if another collection of concepts is used.

children make grammatical errors along the lines of (11a) or (12b), they also tend to attach incorrect meanings to the verb. In fact, not only do the children misinterpret the meaning of the verb, but they tend to assign it precisely the kind of meaning that would give it a semantic structure that would license it to appear in the otherwise ungrammatical form. These results suggest that such sublexical concepts are needed in the psychological study of child language acquisition. They also find their way into theories of adult sentence processing, where, according to some views, the basic semantic structure generated by these sublexical concepts determines how we (initially) parse sentences (e.g., Tanenhaus et al. 1993; Boland 1997). Third, some sublexical concepts even seem to surface in non-linguistically oriented cognitive psychology (e.g., Leslie 1995, 1982; Leslie and Keeble 1987). Leslie presented preverbal infants (some as young as 6 months) with different causal scenarios. The children displayed a differential reaction to the different scenarios, and this reaction tracked the distinction between direct causation and ordinary causation. He then argued that infants possess certain concepts that underlie much of our early understanding of the causal nature of the world. Surprisingly, these concepts are very similar in nature to the sublexical concepts posited by linguists. Finally, in addition to English, linguists frequently employ the same sublexical concepts to theorize about many diverse and typologically unrelated languages, including Japanese, Berber, Chukchee, Icelandic, Turkish, Russian, Tagalog, Malagasy, St'át'imcets, Papago "and many hundreds of other languages" (Hale and Keyser 1993, p. 102; cf. e.g., Comrie 1985; Baker 1988; Tenny 1994; Bittner 1999).

Together, these considerations make it hard to deny that sublexical concepts are part of the design features of human languages in general. I will henceforth assume that the existence of sublexical concepts is not in question.

4. SUBLEXICAL CONCEPTS AND ACCESSIBILITY

We saw above that it would be unsurprising to discover that lexical meanings (like the meaning of *dog*) were accessible. But it is less obvious whether sublexical concepts like CAUSE are accessible or tacit. Since sublexical concepts are one of the more obscure types of semantic features, if they turn out to be accessible, then that will constitute evidence for the general claim that our semantic concepts are accessible. We could explore their accessibility by constructing a test that is structurally analogous to the one we constructed earlier for traces. The sublexical concept I will explore is that of direct causation. Just as we did with traces, we can begin by in-

roducing subjects to the concept by means of some examples, such as the distinction between causative verbs and phrasal causatives. For instance, we might show subjects the data in (14), and explain to them why (14a) means (14b) and not (14c).

- (14) a. Jane boiled the water.
 b. Jane directly caused the water to boil.
 c. Jane caused the water to boil.
 d. Other verbs: *crack, change, tear, close, drain, sink, halt, freeze.*

Speakers can be shown how to pry apart (14a) and (14c) by means of simple thought experiments. For example, suppose Jane told Dave that whatever he does, he is not to boil the water. Suppose also that Dave is a double agent who realizes he can foil Jane's plans by boiling the water, and so he does so. In this circumstance, there is a clear sense in which Jane caused the water to boil; nonetheless it is not true that she boiled the water. It is possible that subjects would need an example like this to distinguish (14a) from (14c). The fact that speakers may need this kind of priming is not a new discovery. Austin's discussion of the dangers of relying on oversimplified models of certain fundamental notions takes causation as a prime example. All too often, Austin observes, we think of causation as having certain prototypical features (e.g., that of a person pushing a stone). The danger with this is that we tend to invalidly import these features into new scenarios that involve some form of causation which may differ from the prototypical cases (Austin 1956, p. 28). Austin showed that in the realm of excuses we may need a much more liberal notion of causation than the one supplied by the oversimple model. In the case of causative verbs, we see that we may need a more *restrictive* one, one that requires that the causation in question be direct, and not mediated, e.g., through the intentions of another agent.

To determine whether subjects have been sufficiently primed for the concept of direct causation, we can see whether they can distinguish causative verbs (e.g., (15a–b)) from their phrasal counterparts (e.g., (15c–d)). Of course, these are only new examples from within the training space, so facility with (15) only indicates a general understanding of the task at hand.

- (15) a. Shanna increased the budget.
 b. Dominic shrunk his shirt.
 c. Shanna caused the budget to increase.
 d. Dominic caused his shirt to shrink.

Just as with traces, the real test concerns how speakers react to other kinds of constructions that are unlike the examples used to explain what direct causation is. A good type of test would involve distinguishing resultative constructions from circumstantials and depictives. Speakers could be given a variety of sentences like those in (16). Just as with traces, their instructions would be to indicate which of the expressions contain direct causation as a part of their meaning, with special attention being paid to the underlined parts.

- (16) a. Jenny ate the meat raw.
 b. Alison whipped the cake batter smooth.
 c. David wrote a term paper drunk.

Here (16a) simply means that Jenny ate the meat while the meat was raw, and (16c) means that David wrote a term paper while he was drunk. But (16b) is a resultative construction: it means that Alison whipped the cake batter with the directly caused result being that the batter became smooth. (Strictly speaking, these sentences are ambiguous, so they might need to be embedded in a disambiguating context.) It seems highly likely that once they are suitably primed, speakers will be quite good at spotting the resultative constructions and distinguishing them from the circumstantials and the depictives. In any case, they are sure to be much better at this task than they were at spotting traces.⁹

The fact that speakers should be able to identify sublexical concepts like direct causation in novel environments, although they cannot identify traces in novel environments suggests that there is a real epistemic difference between syntactic and semantic features. The examples of traces and direct causation lend particular credence to the claim when we note just how simple and basic traces are, relative to many other syntactic features, and just how obscure and unapparent sublexical concepts are, relative to many other semantic features. This sort of evidence supplies a substantial amount of evidence for what I will call the *semantic accessibility thesis* (or SAT for short). According to SAT, the semantic features of a sentence are accessible to speakers of the relevant language. Thus stated, there are several ways to interpret SAT. Let me now fill out what I take to be the most interesting version of SAT. In the first place, a bit more should be said about the semantic features of a sentence. Although I stipulated early on that the view of language I am working with is internalist (in the sense of Chomsky's I-language), there is still another issue: Should the semantic

⁹ I have confirmed the empirical claims about accessibility, and the results are reported in Johnson (2002).

features in question be those that are actually part of the speaker's language (i.e., actually part of that state of her mind that enables her to acquire and speak the language she does), or should they be those features that the speaker takes her language to have? (I am grateful to Jim Higginbotham for drawing my attention to this question.) After all, one might easily be mistaken in thinking that one's language has (or lacks) some particular feature. It is unclear if there is a nontrivial version of SAT which says that speakers have access to the semantic features of sentences that they take them to have. In any case, the evidence we have seen supports the stronger version of SAT which says typical speakers can access the semantic features of sentences that their linguistic abilities actually endow them with. I'll have a bit more to say about this aspect of SAT in the next section, but for now I will focus on specifying the content of SAT. The second way that SAT is underspecified concerns what counts as being aware of a semantic feature. Here too, a question arises: Does SAT predict that speakers can access the semantic impact that a given semantic feature has on a sentence in which it is contained, or does it make the stronger prediction that speakers can access the semantic feature as such? That is, when speakers access a semantic feature, do they access it *qua* semantic feature of the sentence, or do they only access it in the sense that they understand the sentence well enough to exhibit a grasp of the part of the meaning (or truth-conditions) of the sentence centered around that feature? Here again, the evidence we have seen so far supports the stronger view of SAT, according to which we can access the semantic features as such. (We will see in the next section that this option is a kind of upper bound on the possible strength of SAT.) With that said, we can specify the thesis as follows:

(SAT) *Semantic Accessibility Thesis*: Ceteris paribus, the actual semantic features of sentences of a typical speaker's language are accessible as such by that speaker.

Of course, there is much work left to be done in terms of confirming SAT. Moreover, specifying the ceteris paribus constraints under which such a claim holds may well be nontrivial. Nonetheless, the evidence just considered constitutes an important first step. Moreover, I think that sublexical concepts present the most difficult challenge to SAT. I turn now to some further discussion of the thesis.

5. OBJECTIONS, REPLIES AND DEVELOPMENTS

In this section, I will fill out some of the details of SAT by addressing a number of issues and worries that arise for the thesis. The first issue I wish to address concerns some of the compositional elements of language. One might resist SAT because some expressions are standardly given their semantics in terms of a logical apparatus that exploits set theoretic and model-theoretic vocabulary. Assuming that these semantic theories are on the right track, the objection goes, SAT incorrectly predicts that typical speakers can access a good deal of mathematical structure. To see how this objection works, consider the example of a quantifier like *all* or *no*. For various reasons, quantifiers are standardly taken to express higher-order relations between properties. Leaving aside many details, it is common for semantic theories to express that e.g., *All men are happy* is true if and only if $\{x: x \text{ is a man}\} \subseteq \{x: x \text{ is happy}\}$, and *no men are happy* is true iff $\{x: x \text{ is a man}\} \cap \{x: x \text{ is happy}\} = \emptyset$. But now it looks like in order to access the meaning of our sentences, we have to be able to intuit, with only a little priming, various facts about the algebra of sets. As more details are added to the semantic theory, this mathematical apparatus becomes more and more complicated (cf. e.g., Bittner 1999 for an example of this complexity). This apparatus is surely not accessible to untrained speakers of the language.

My reply to this worry is that it appears to confuse nature and its representation. On the one hand, there is the semantic structure of human languages, which is what a formal semantic theory represents. On the other hand, there is the mathematical structure a semantic theory uses to represent the semantic structure. There is simply no reason to believe that all of the mathematical apparatus used in the representation of semantic structure is a straightforward component of semantic structure. In the typical psychological case, this mistake is never made. For instance, suppose that a theory of how visual search in some particular domain works predicts that the number of milliseconds the subject will take to indicate where the target object is (Y) a linear function of the number of non-target items in the display (X), say $Y = 40X + 600$. Even if we assume that subjects have some kind of tacit awareness of their visual search strategy, we need not suppose that we are committed to subjects being tacitly competent at algebra. In such a situation, we would look for some other interpretation of the mathematical data. In the present case, we might hypothesize that it takes 600 ms to implement the decision (such as tapping the appropriate key on the keyboard), and 40 ms to examine each item using a random search strategy with only one target item in the display. The tacit aware-

ness represented by this function would probably not be hypothesized to be awareness of the function itself, but rather awareness of the relevant components of a search procedure, and perhaps the time each component requires for execution.) There seems to be no reason for doing otherwise in the case of formal semantics. Rather than assuming that every detail of the mathematical apparatus of a semantic theory is “psychologically real” in some interesting sense, its purpose is simply to clarify and perhaps organize certain aspects of the empirical phenomenon in question.

Another way to observe this point is to note that simply because formal semantic theories use a mathematical structure to represent the structure present in the mind does not imply that the particular structure used has more psychological reality than some other structure that might have been used. For instance, simply because a linguistic theory uses, e.g., an algebra whose domain is generated from a base of primitive meanings closed under a class of semantic composition functions, that does not make the details of that mathematical tool for representing meaning more psychologically relevant than, e.g., a proof-theoretic system where axioms replace primitive meanings, and inference rules replace the functions. Since either theory could have been used just as well, the mathematical differences between them are surely not psychologically real. Moreover, there are infinitely many more formal theories that are identical in terms of their empirical predictions, and so none of the formal features peculiar to less than all of these theories can be said to be linguistically real. The mathematical commonality between all these theories is bound to be very general, and it seems quite plausible that typical speakers will have that much grasp of the semantic structure of the language. A mathematical system that captured all *and only* the relevant properties of human languages would command more attention than its contemporary competitors as informing us about the true nature of combinatorial semantic structure.¹⁰ However, barring the ad hoc systems that have been created, no such system has yet been found. And if one ever is, there is no reason at present to suppose that it will present any problems for the current argument.

In short, like Quine’s distinction between a grammar’s “guiding” or “fitting” a speaker’s behavior, we can expect there to be infinitely many different mathematical structures, all of which correctly represent the formal structure of language (and all of which contain a great deal more structure as well) (Quine 1972). But unlike Quine, we need not claim that psychological reality is only responsive to the extension of the grammar. As Evans

¹⁰ In fact, the theory need only capture many of the properties of some interesting sub-domain of language, in such a way that the mathematical “joints” of the theory correlated in some interesting way with the apparent major structural features of language.

(1981) and others have noted, many factors may support empirical claims about the superiority of one theory of language over another extensionally equivalent one. Clearly it is possible to attribute a great deal of formal structure to human languages without attributing to them all the formal structure present in the model one happens to use. All of this is consistent with SAT.

A second worry that one might have with SAT concerns the fact that there are disagreements in linguistics and the philosophy of language. The mere fact that two speakers of the same language can disagree about the meaning of an expression suggests that they do not both have access to the meaning of an expression. But professional linguists and philosophers of language frequently disagree about the semantic structure of various expressions. Furthermore, they are surely amply primed to the relevant issue, so they should be well-suited to make occurrent any accessible beliefs that they have about language. So if speakers have access to the meaning of proper names, for instance, then there should not be any disagreement among professionals about whether they are (semantically speaking) descriptions or directly referring terms. The fact that the semantic structure of proper names is a topic of vigorous debate seems to almost guarantee that we cannot access the meaning of proper names. More generally, the fact that trained professional linguists and philosophers of language often produce false semantic theories seems to show that we cannot access the meaning of parts of the language (cf. Dennett 1978, p. 304 for discussion).

The above worry is fueled by SAT's apparent denial of the claim that typical speakers cannot simply intuit (even when given a bit of priming) whether proper names are devices of direct reference or disguised definite descriptions (or something else). Although there is more to be said about this claim, I will here confine myself to the observation that SAT is neutral with respect to it. SAT states that speakers can access the semantic features of sentences and use these features in overt rational cognitive activity. SAT does not additionally assert that typical speakers can recognize that a given structure is the correct semantic structure of an expression. To have this latter ability, typical speakers would need to have (i) an ability to form certain beliefs about the semantic structure induced by these features, and (ii) an ability to reliably attribute greater plausibility to that belief than to any other beliefs about the semantic structure of the expression. (ii) is no part of SAT, and it is because we lack (ii) that linguistics and the philosophy of language are difficult areas of research. To see what SAT says, recall that the primary goal is to contrast our awareness of the semantic properties of a language with our awareness of its syntactic properties. SAT is a generalization of the marked difference in graspability between these two:

our ability to identify and use semantic features in overt reasoning tasks is better and faster than our ability to do the same with syntactic features. It is in this sense that the semantic features of natural language appear to be accessible to speakers, and it is also in this sense that our grasp of semantic features is not tacit. But simply because the semantic features of language are accessible to us in this sense, nothing follows about the strength of a belief of the form “Expression E has semantic structure S” to override other conflicting beliefs, even if the former is occurrently entertained. In short, although SAT predicts that we can be non-tacitly aware *of* certain semantic features, and that we can be non-tacitly aware *as* semantic features, it does not predict that we will always be non-tacitly aware *that* they are the semantic features which form the true semantic structure of the expression. This is the upper bound I mentioned before. Although there is a strong and interesting sense in which we can access the semantic features of our language, there are also epistemic limits to our awareness.

At this point, it may be useful to consider some of the empirical bite SAT has. SAT predicts that sublexical concepts are accessible (a claim which conflicts with what we find in, e.g., Pinker 1989, p. 359). There is no *a priori* reason why such a claim should be true. It is certainly conceivable that our linguistic abilities would be such that a true theory of the structural level of semantics would posit an array of sublexical concepts which had such bizarre forms that they were simply unusable and unrecognizable by speakers. It could have turned out, for instance, that the most fundamental concepts would be ones which applied to such a motley bunch of things that we would not be able to describe the concept in English at all. Since much of our reason for positing sublexical concepts in the first place had to do with the distribution of lexical features, it is possible that the distribution of these features would have been best accounted for by positing a bunch of intuitively bizarre concepts. (Consider, e.g., Perlmutter’s original classification of unaccusative verbs (Perlmutter 1978), and suppose a theory implied that this classification was captured by a single primitive sublexical concept.) We can also suppose that data from cross-linguistic, psycholinguistic, and psychological studies would confirm that the sublexical concepts structuring our grammar are these bizarre ones. In such a case, not all semantic features would be accessible, and SAT would be false. However, a look at the sorts of sublexical concepts actually posited by linguists and psycholinguists does not support this possibility.

Although SAT has not been discussed in the literature, some researchers have presented views that would suggest that it may be of limited importance. For instance, Higginbotham writes,

The sharp distinction between what is available, either off the bat or upon reflection, to us as guiding our behaviour and what is for ever sealed off from consciousness, so that we can no more be aware of ourselves as following the principles in question than we can be aware of the ways in which visual information gives us the boundaries of objects, has no evident role to play in the construction of linguistic theories. (Higginbotham 1998b, p. 440)

However, it is not clear that this is so. If SAT is true, it constrains linguistic theories so that they do not employ inaccessible semantic features. Such an empirical constraint addresses a problem that Fodor has raised for theories of lexical semantics (e.g., Fodor 1998, chap. 3). Fodor's worry is that it is all too easy to find concepts that correctly organize words into appropriate grammatical classes. Once we have distinguished various classes of words by their grammatical properties, for each such class, we can simply construct a semantic property that is common and unique to the words in the class, using whatever Goodman-esque techniques we wish. In reply to this, though, we can note that SAT prevents just any old "gruesome" semantic feature from being employed, by requiring accessibility of all of them. Thus, finding semantic properties that correspond to distributionally identified word-classes becomes that much harder of a project, and the successes of research into this area of linguistics becomes that much more compelling.

In addition to its ability to constrain linguistic theories, SAT also raises some interesting questions about the respects in which, and the extent to which, our linguistic abilities are modular. Consider the case of sublexical concepts. Some linguists and psycholinguists have suggested that sublexical concepts are not restricted to the language processing mechanism (e.g., Pinker 1989, p. 359 and Grimshaw 1994, 1994b). They hypothesize that the process of lexical acquisition begins with the child pairing a word with a concept (e.g., the word *spray* with the concept of spraying). Once the child has done this, the language mechanism inspects the concept and records the general structure of the concept, using the vocabulary of a tightly restricted language of sublexical concepts. According to this view, the sublexical concepts are used only in the language mechanism, and the other "cognitive" semantic properties of expressions are delivered from language processing into higher cognition (or vice-versa). (I don't mean to suggest that this is anybody's entire view of lexical acquisition; cf. e.g., Gleitman (1990) for some potentially relevant additional considerations.) But a puzzle about sublexical concepts arises for such a view. If sublexical concepts are restricted in their application to the language processing mechanism, then why do they all appear to be accessible? It's easy to believe, for instance, that many syntactic features of language are restricted in this way, since speakers show little awareness of them. But if sublexical concepts are similarly restricted, it is a little surprising that we

are so readily aware of them. Of course, there are several replies that a proponent of the view that sublexical concepts are restricted to the language mechanism might make. One possibility is that because these sublexical concepts play such an active role in the linguistic aspects of our rational lives, we are readily prepared to acquire new concepts – ones that can be used in ratiocination, not ones that are components of language – when the need arises. Thus, it may be that upon being primed with information about what kind of concept is relevant, speakers are able to quickly construct a concept that matches the sublexical concept in question. While some suggestion along these lines may work, there would still remain the question why such an explanation should hold for semantic features like sublexical concepts, but not for syntactic features. After all, the relevant syntactic features appear to play an equally active role in the linguistic aspects of our rational lives. Or if they don't play such a role, one would like an explanation of what precisely is the difference between syntactic features and semantic features like sublexical concepts. A full-blown discussion of the difference between syntax and semantics, and the various roles the two play in our cognitive lives is beyond the scope of this paper. So, having noted this issue and the role that our investigation can play in it, I will move on.

Finally, notice that we can answer a question that was raised back in section 2. There we asked whether the inaccessibility of linguistic features such as traces might be due to the fact that they are theoretical entities posited not because they are immediately perceivable but because they are an important part of a total theory. Such a possible explanation is undermined by the fact that sublexical concepts are theoretical entities in the same sense that traces are, yet they are accessible. So a successful explanation of why syntactic features like traces are inaccessible will need to appeal to something more than their theoretical status.

6. PHILOSOPHICAL APPLICATIONS

So far, much of our discussion has concerned some relatively detailed and narrow issues. In this section, I will explore how the discussion fits into two broader philosophical issues. The first issue concerns whether we think in the same language we speak, and the second issue concerns how to explain our knowledge of our public language.

6.1. *Do we Think in English?*

There is a longstanding philosophical debate concerning whether we think in English (or in our first language, more generally). Suppose, for example, you hear a sentence uttered and you understand it and even come to believe it. Is the syntax of the sentence you heard part of your belief? More generally, is thinking a way of “speaking silently to yourself”? Many philosophers have defended this claim (e.g., Harman 1973, 1975, p. 271, 1999, ch. 9–10; Sellars 1956, 1969) or at least argued against its implausibility (e.g., Ludlow 1999, pp. 22–26, 165–169). In contrast to these arguments, though, our discussion of the (in)accessibility of various linguistic features seems to support a general argument that we do not think in English. The basic idea behind the argument is that there isn’t much reason to think that the complex details of natural language syntax, which partly constitute the language, are part of the corresponding thoughts, in any interesting sense. The form of this argument is as follows:

- (17) A typical speaker cannot access many of the syntactic properties of her natural language.
- (18) *Ceteris paribus*, if a typical speaker cannot access a feature of her language, then that is evidence that the feature is not a part of the corresponding thought.
- (19) Thus, there is evidence that much syntactic structure of language is not part of speakers’ corresponding thoughts.
- (20) We think in our natural language only if our thoughts have the syntax of that language.
- (21) Thus, there is evidence that we do not think in our natural language.

I endorse this argument. The inferences in it are valid, and I think that each of its premises are plausible. (17) was defended in section 2, and I take (20) to be unproblematically true. The only remaining premise is (18), which says that if you can’t tell whether something is a part of your thought, assume it is not unless you have good reason for doing otherwise. Thus (18) appears to be an instance of Occam’s Razor: posit no more structure in thought than is necessary. So it would seem that (18) is secure. However, in discussing this argument, I have encountered a number of objections to it, all of which involve arguing that the *ceteris paribus* clause of (18) goes

unsatisfied. If this is right, then my use of (18) above is illegitimate. I will spend the rest of this section canvassing these objections.

The first objection I have encountered was offered independently by several linguists in conversations where SAT had not been introduced. According to this objection, when you entertain the meaning of a perceived sentence of your language, we may plausibly suppose that the semantic properties of the sentence enter into, or partly constitute, our thoughts. Since some of these properties enter into our thoughts, the simplest hypothesis is that all semantic properties do so. (In linguistic terms: perhaps our linguistic abilities interface with the conceptual-intentional system in such a way that all semantic features are a part of the latter, and our linguistic abilities are merely a set of instructions as to how to assemble these features into a thought.) But, the argument goes, surely there are semantic features which are inaccessible to us. These inaccessible features are, by hypothesis, part of the thought. But if some inaccessible linguistic features partly constitute our thoughts, the objection concludes, we do not yet have a reason for supposing that syntactic features are not part of our thoughts, too. Unsurprisingly, I am unpersuaded by this objection. If SAT is true, then one of the premises of the objection is false. SAT suggests that all semantic features are accessible, so we cannot slip inaccessible features into thought by this route. So by accepting SAT, we may continue to doubt that natural language syntax plays any interesting role in thought in virtue of its being inaccessible.

Another objection I have heard begins by conceding that syntax and semantics are cognitively different, and yet still maintains that (18) does not apply to syntax. According to this objection, the syntax of sentences needs to be present as (part of) the object that is interpreted. One shouldn't be surprised that the syntax is inaccessible, because it is the vehicle of interpretation, not the interpretation itself. But if syntax is what the interpretation is structured on, much like a skeleton structures the visible skin and muscles of a body, then we do have reason for supposing that syntax is part of our thoughts. In reply to this, though, we can note that in order to defend a claim about the proper cognitive vehicles of thought, we need to know more about the interaction between language processing and thinking than we currently do. Moreover, the claim that we think in English does not amount to the claim that logically perspicuous formulae (with only the syntax necessary to support such formulae) enter into one's thinking. For almost all syntactic frameworks, any level of syntax contains a good deal of information that is both inaccessible and semantically irrelevant.¹¹ If we

¹¹ There are hypotheses that appear to belie this last claim. A strong form of Minimalism would hold that the principle of Full Interpretation requires that there be no semantically

think in English, then presumably even our nonverbal thoughts would have an English syntax, too, including those features that are generated solely for linguistic reasons. So when you were thinking about what you had for breakfast this morning, there were various syntactic structures present in your mind (such as devices involved in case-checking) whose existence is justified because of the work it does in accounting for how typical speakers do and do not speak. Such an outcome is of course theoretically possible. But without any reasons for supposing that all this structure is present in thought, this kind of view increases the complexity of the structure of thought simply in order to maintain a theory, in direct violation of Occam's Razor.

Another way to see this last point is to contrast what is known about the structure of natural language with what is known about the structure of propositional thought. On the one hand, there is quite a bit that is known about the structure of natural language sentences. This is so because natural language admits of phonological (and perhaps orthographic) realization, which enables us to isolate and study the outputs of a human grammar more or less directly. Moreover, we can often distinguish sentences from similar objects that are not the outputs of the grammar. The ability to isolate sentences from non-sentences is a primary source of evidence for positing phonologically and/or semantically null features of natural language. When it comes to thoughts, however, the situation is different. We do not have the same ability to individuate, isolate, and study thoughts with anything like the ease or exactitude of our abilities to study sentences. Similarly, we can't distinguish thoughts from non-thoughts with anything like the degree of fine-grainedness of our grammaticality judgments. What would it be, for instance, to contrast a very similar pair of putative mental entities and come up with the judgment that one of them is a thought but the other is not? Of course, my occurrent belief about what I had for break-

irrelevant material at LF, the level of syntactic description at which semantic interpretation takes place. (More precisely, that all the material in an LF structure is relevant to the conceptual-intentional system with which the LF structure interfaces) (cf. Chomsky 1995; Marantz 1995; the philosophical foundations of Minimalism are discussed in, e.g., Chomsky 2000, ch. 1, 1995, ch. 3). Although I cannot argue for it in detail here, I will suggest that in order to make this claim relevant to present considerations, such a theory would also have to claim that the vast amounts of morphological structure that is employed in lieu of additional syntax must also be accessible. The resulting claim is an extremely strong one, for which there is little or no evidence. In this respect, I am happy to adopt a wait and see attitude. (If such an extremely strong form of Minimalism is correct, then it may be that my theory will turn out to *support* the claim that we think in English, by showing that all the (*ex hypothesi* semantically relevant) parts of Minimalist syntax actually denote concepts that are accessible by us.) I am indebted here to discussion with Paul Pietroski.

fast is a thought and the neural aspects of my current chewing behavior is not a thought, but we don't generate contrasts anything like the linguistic judgments that *It is John that left* is a sentence but **It is John left* is not. Our inability to produce independent evidence about the nature of thought means that we do not have the kind of reasons for positing unperceived structure in thoughts that we have for positing such structure in sentences. So if we don't perceive structure in thought, then (unlike natural language) unless we find some other justification for positing that structure in our thoughts, we don't have good reason to hold that that structure is there anyway. And this is what (18) says.

There are two more objections to my use of (18) that I will briefly discuss. The first one states that language and thought appear to be relevantly similar, not least in that both have semantic properties. Thus, since the former has a certain syntax, we might treat this as evidence that the latter has this syntax, too. I will call this the *argument by analogy*. The other objection is that language and thought are also similar in that they seem to always co-occur: entities seem to have language iff they have thought (at least the sort of complex propositional thought distinctively engaged in by humans). Perhaps this co-occurrence can be explained by supposing that language and thought are really one and the same thing. If so, then it would appear that we do think in our natural language. Call this the *co-occurrence argument*. In reply, both arguments appear to be undermined by the empirical facts about language. The argument by analogy is certainly right to point out that there are some deep and important affinities between language and thought. But there are also many deep differences. For instance, all natural languages have phonological properties. Surely, however, there is no reason to suppose this is true of thought. Relatedly, many linguists maintain that the phonological system of a language may have grammatical effects. "Heavy NP-shift" is allowed when the relevant features have the right phonological "weights", as in *I kicked against the wall every ball I could find that was not too flat to bounce*. However, they argue, when the phonological weights are reversed, the resulting sentence is no longer grammatical, as in **I kicked against the wall of the room where I was born and spent my youth dreaming about my life the cat*. I know of no evidence that thoughts have phonological properties, or that these properties can affect whether a given structure is a thought or not. Furthermore, language and thought may also differ in an entirely different respect. As Higginbotham has emphasized (e.g., 1985, p. 3, 1989, p. 154), our ability to understand language entails an ability to understand certain strings that are not sentences of the language. Although it is ungrammatical, **The child seems sleeping* is fully interpretable. If we think in English it is hard to see

why we can understand this string at all. (Of course, we could posit an ability to translate it into an English sentence, but would we do this for any other reason than to maintain that we think in English?)

The problem with the co-occurrence argument is even more straightforward. This argument suggests that language and thought are the same. But there are good reasons to suppose that language processing and higher cognition are by and large located in different regions of the brain and realized by distinct kinds of neurological activities. People with certain kinds of disorders such as Williams' syndrome may have severely impaired reasoning abilities, even though their linguistic abilities remain intact (e.g., Bellugi et al. 1999; Pinker 1999; Deacon 1997, pp. 264–278; Ullman et al. 1997). In the other direction, people with aphasia may exhibit relatively normal abilities for reasoning even though their language abilities are severely impaired (e.g., Stromswold 1999; Ullman et al. 1997; cf. Pinker 1999, ch. 9, for informal discussion and many citations). Of course, much of the data regarding the linguistic and cognitive abilities of subjects with various kinds of neural disorders is sketchy and inconclusive. Nonetheless, the large body of evidence from imperfectly functioning brains does point toward the hypothesis that linguistic and reasoning abilities can be prized apart, even if this rarely happens in normal cases.

In sum, there is little reason to suppose that we think in English, primarily because there is little reason to think that any useful characterization of thought will entail that English syntax is a part of the structure of our thoughts. Notice, though, that the argument leaves it open that our thoughts somehow involve the semantics of natural language. While the semantics of one's language surely has some important connections to our ability to think, the exact connections between, e.g., our word-meanings and our concepts, and our sentence-meanings and our complete thoughts, is the topic for another project.

6.2. *Knowledge of a Public Language*

The second philosophical issue I will consider concerns the explanation of how we know, and are presumed by other speakers to know, the meanings of our public language. We can see the epistemological problem by distinguishing a speaker's public language from her individual language or "I-language" (cf. Chomsky 1986). By hypothesis, we may suppose that a speaker understands her own individual language, in virtue of the fact that she (tacitly) knows a grammar of that language. However, public languages are social phenomena, and their structure needn't be determined by the individual language of a particular speaker. Nonetheless, my judgments about the meanings of English sentences (construed as a public language)

are more authoritative than would be the judgments of a monolingual speaker of Chinese. What is it about me and my individual language that makes my beliefs about my public language justified? Epistemological problems of knowledge of one's public language have a long history, going back to the seminal works of Quine, Dummett, and Davidson (e.g., Quine 1960; Dummett 1975, 1976; Davidson 1973, 1974; some more recent studies include George 1990; Smith 1998; Higginbotham 1989, 1998, 1998b).

I want to suggest that we can get a partial grip on this epistemological issue by appealing to the nature of sublexical concepts. We have seen evidence that the semantic structure of simple sentences of speakers' individual languages is due in large part to the sublexical concepts present in the words. There is, I think, good reason to suppose that this structure is present in expressions of the public language, and that speakers believe this with good reason. In what follows, I will briefly sketch out the argument I have in mind. In the first place, to see that this structure is present in the public language, note that the latter's structure supervenes in some relevant sense on the structures of the individual languages of the members of the relevant linguistic community. This means that a public language will have a structural property if the individual languages of the members of the linguistic community have it too. But individual languages in general have the property of being such that their simple clauses are structured by the semantic properties associated with sublexical concepts. Moreover, we have seen evidence that this structure is relatively invariant across speakers, and that it is productive over time, since it correctly predicts the structural properties of new words as they enter an already existing language. The ubiquity of sublexical semantic structure can be contrasted with the varying ways that the members of a population may conceive of the meaning of a word in their public language. Many speakers, for example, conceive of the meaning of the word *livid* as meaning red – to be livid is to be red with rage – while other speakers and most dictionaries take it to mean pale (this example has been discussed by Chomsky). In this case, it is simply not obvious what the public language meaning of *livid* is, or what principles of language use would determine an answer. (In this case, I suspect that there is no answer if we individuate public languages along intuitive lines.) However, the same is not the case for sublexical concepts. Nobody misinterprets *Bill shot John dead* as meaning Bill attempted to shoot John dead, or as meaning Bill shot John because John was dead (although they could interpret it as a resultative, circumstantial, or depictive cf. (16)). The semantic structure of these constructions is fixed throughout the population of speakers. Thus, unlike many other components of meaning, we may take

the semantic structure that is given by sublexical concepts as a component of the public language.

We have just drawn a metaphysical conclusion about (part of) the structure of a public language by appealing to the fact that speakers of that language are in strong collective agreement about the way its sublexical concepts serve to structure the language. We can also use this fact to support an epistemological claim about a speaker's beliefs about the general meaning of the simple sentences of her language. The argument goes like this. A given speaker *S* has an accessible awareness of the semantic structure of the clauses of her individual language, and she behaves accordingly. Moreover, barring evidence to the contrary, *S* will believe that her public language has the features of her individual language. *Ceteris paribus*, *S* is *not* justified in supposing that an expression of her public language has the same semantic structure of the corresponding expression of her individual language only if other speakers provide evidence that it is not. But the only way that other speakers would do this is by behaving as if the expression was not so structured. But as we saw above, speakers are in broad agreement about this semantic structure, so they too will behave in a manner consistent with *S*'s hypothesis. Thus, concerning expressions of her language, *S* will have justified beliefs about their meaning, at least insofar as this meaning is determined by sublexical concepts. A similar conclusion has been drawn by Higginbotham:

If it were always possible of one's conception of the meaning to deviate from the meaning [sc. of the expression of the public language], then there would be no entitlement to knowledge of one's own [sc. public] language at all, contrary to fact. But there is something left over when the normative is deleted: namely, the realm within which meaning is what it is because of the way we are made – that is, because of the (for us) inevitable structure of the subpersonal system. (Higginbotham 1998b, p. 438)

I will conclude this discussion by saying briefly why this kind of semantic knowledge (or entitlement) is important. In order to best exploit what Austin calls "natural economy of language" (Austin 1956, pp. 16, 21), speakers must grasp facts about the semantic structure of clauses. For instance, the direct object of verbs must be totally "affected" by the action of the verb (cf. Tenny 1994). The sentence *John loaded the boxes onto the wagons* is true only if all the boxes were loaded onto the wagons, even though not all the wagons need to have been loaded with boxes. In contrast, the truth of *John loaded the wagons with the boxes* requires that all the wagons were loaded, even though not all the boxes need have been loaded onto the wagons. It is your knowledge of the general architectural features of your public language that enables you to recognize this fact and to use it in communication. If members of a linguistic community did not

share this understanding about the semantic import of direct objects, they would experience fundamental difficulties in communication. (A failure to understand this restriction could lead to disasters with just about any transitive verb. Suppose I am unaware of this constraint on direct objects, and I tell you that I have killed the alligators that were swimming in the pond, believing to have done so since I killed one such alligator.) On the other hand, when members of the same community have a shared understanding of the semantic structure provided by sublexical concepts, a certain amount of communication can be effected even if not all of the vocabulary is shared. For instance, if I tell you *Mary skinked the dishes onto the table*, you will be able to tell that I am describing something that happened to the dishes, even though you do not understand the verb *skink*. Moreover, you will also be able to discern that some portions of the table may not have been affected by the skinking. You can extract the general structure of this clause by means of your understanding of the role sublexical concepts play in structuring the grammar of the clause (for related remarks about syntax, cf. Higginbotham 1989b). Thus, the mere ability to extract the structure provided by sublexical concepts enables a hearer to determine the general kind of event that is being reported.

7. CONCLUSION

We have seen that our epistemic relationship to language is anything but simple. Our awareness of the syntactic features of language appears to be by and large tacit, whereas our grasp of semantic features appears to be accessible. At the same time, our awareness of semantic features is not so strong as to deliver beliefs about semantic structure of the sort that would make linguistics and philosophy of language easy. The accessibility of semantic features, in particular sublexical concepts, we saw, has various linguistic, cognitive, and philosophical implications.

Sublexical concepts are interesting not least because they are one of the only things that interact heavily with grammar, linguistic meaning, and higher cognition. But there is much room and need for further philosophical attention to sublexical semantics and the nature of the lexicon. Such questions as what sublexical concepts there are and to what extent they are accessible leap to mind. There is also work to be done in determining the interesting boundaries between syntax and semantics. I have offered some evidence that there is an epistemic difference between the two. However, as discussed in section 1, the accessible/tacit distinction is vague. It may be that the distinction between syntax and semantics is also vague. (Personally, I highly doubt that there is any useful sharp distinction between syntax and semantics.) In such a case, the nature and import of SAT may

change: perhaps our grasp of properties of the sentence will be vague in exactly the way syntax is vaguely distinct from semantics, perhaps not. All this is, of course, a matter for further research.

REFERENCES

- Austin, J. L.: 1956, 'A Plea for Excuses', *Proceedings of the Aristotelian Society* **LVII**, 1–29.
- Baker, Mark: 1988, *Incorporation*, University of Chicago Press, Chicago.
- Bellugi, Ursula, Debra Mills, Terry Jernigan, Greg Hickok, and Albert Galaburda: 1999, 'Linking Cognition, Brain Structure, and Brain Function in Williams Syndrome', in H. Tager-Flusberg (ed.), *Neurodevelopmental Disorders*, MIT Press, Cambridge, MA, pp. 111–136.
- Bittner, Maria: 1999, 'Concealed Causatives', *Natural Language Semantics* (7), 1–78.
- Boland, Julie E.: 1997, 'The Relationship between Syntactic and Semantic Processes in Sentence Comprehension', *Language and Cognitive Processes* **12**(4), 423–484.
- Burge, Tyler: 1990, 'Frege on Sense and Linguistic Meaning', in D. Bell and N. Cooper (eds.), *The Analytic Tradition*, Blackwell, Oxford, pp. 30–60.
- Chomsky, Noam: 1965, *Aspects of the Theory of Syntax*, MIT Press, Cambridge, MA.
- Chomsky, Noam: 1981, *Lectures on Government and Binding*, Foris, Dordrecht.
- Chomsky, Noam: 1986, *Knowledge of Language*, Praeger, Westport, CT.
- Chomsky, Noam: 1995, *The Minimalist Program*, MIT Press, Cambridge, MA.
- Chomsky, Noam: 2000, *New Horizons in the Study of Language and Mind*, Cambridge University Press, Cambridge.
- Comrie, Bernard: 1985, 'Causative Verb Formation and Other Verb-Deriving Morphology', in T. Shopen (ed.) *Language Typology and Syntactic Description*, Vol. III, Cambridge University Press, Cambridge, pp. 309–348.
- Crimmins, Mark: 1991, 'Tacitness and Virtual Beliefs', *Mind and Language* **7**(3), 240–263.
- Culicover, Peter: 1997, *Principles and Parameters*, Oxford University Press, Oxford.
- Davidson, Donald: 1973, 'Radical Interpretation', reprinted in Davidson 1984, pp. 125–139.
- Davidson, Donald: 1974, 'Belief and the Basis of Meaning', reprinted in Davidson 1984, pp. 141–154.
- Davies, Martin: 1981, *Meaning, Quantification, Necessity*, Routledge & Kegan Paul, London.
- Davies, Martin: 1989, 'Tacit Knowledge and Subdoxastic States', in A. George (ed.), *Reflections on Chomsky*, Blackwell, Oxford, pp. 131–152.
- Deacon, Terrence W.: 1997, *The Symbolic Species*, W. W. Norton, New York.
- Dennett, Daniel: 1978, *Brainstorms*, Bradford Books, Cambridge, MA.
- Dowty, David: 1979, *Word Meaning and Montague Grammar*, D. Reidel, Dordrecht.
- Dummett, Michael: 1975, 'What is a Theory of Meaning? (I)', in *The Seas of Language*, Clarendon Press, Oxford, pp. 1–33.
- Evans, Gareth: 1981, 'Semantic Theory and Tacit Knowledge', reprinted in Evans (1985), pp. 322–342.
- Evans, Gareth: 1985, *Collected Papers*, Clarendon Press, Oxford.
- Featherston, Samuel, Matthias Gross, Thomas F Münte, and Harald Clahsen: 2000, 'Brain Potentials in the Processing of Complex Sentences: An ERP Study of Control and Raising Constructions', *Journal of Psycholinguistic Research* **29**(2), 141–154.

- Fodor, Jerry A.: 1968, 'The Appeal to Tacit Knowledge in Psychological Explanation', *The Journal of Philosophy* **65**(20), 627–640.
- Fodor, J. A.: 1975, *The Language of Thought*, Harvard University Press, Cambridge, MA.
- Fodor, Jerry A.: 1990, Review of Stephen Schiffer's *Remnants of Meaning*, in *A Theory of Content and Other Essays*, MIT Press, Cambridge, MA.
- Fodor, Jerry A.: 1998, *Concepts*, Clarendon Press, Oxford.
- George, Alexander: 1990, 'Whose Language is it Anyway? Some Notes on Idiolects', *The Philosophical Quarterly* **40**(60), 275–298.
- Gleitman, Lila R.: 1990, 'The Structural Sources of Verb Meanings', *Language Acquisition* **1**(1), 3–55.
- Graves, Christina, Jerrold J. Katz, Yuji Nishiyama, Scott Soames, Robert Stecker, and Peter Tovey: 1973, 'Tacit Knowledge', *The Journal of Philosophy* **70**(11), 318–330.
- Grimshaw, Jane: 1994, untitled manuscript, Rutgers University.
- Grimshaw, Jane: 1994b, 'Lexical Reconciliation', *Lingua* **92**, 411–430.
- Gropen, Jess, Steven Pinker, Michelle Hollander, and Richard Goldberg: 1991, 'Syntax and Semantics in the Acquisition of Locative Verbs', *Journal of Child Language* **18**, 115–151.
- Haegeman, Liliane: 1994, *Introduction to Government and Binding Theory*, MIT Press, Cambridge, MA.
- Hale, Kenneth and Samuel Jay Keyser: 1987, 'A View from the Middle', Lexicon Project Working Papers 7, Center for Cognitive Science, MIT, Cambridge, MA.
- Hale, Kenneth and Samuel Jay Keyser: 1993, 'On Argument Structure and the Lexical Expression of Syntactic Relations', in K. Hale and S. Keyser (eds.), *The View from Building 20*, MIT Press, Cambridge, MA, pp. 53–109.
- Harman, Gilbert: 1973, *Thought*, Princeton University Press, Princeton.
- Harman, Gilbert: 1975, 'Language, Thought and Communication', in Keith Gunderson (ed.) *Minnesota Studies in the Philosophy of Science*, University of Minnesota Press, Minneapolis.
- Harman, Gilbert: 1999, *Reasoning, Meaning and Mind*, Clarendon Press, Oxford.
- Higginbotham, James: 1985, 'On Semantics', in E. Lepore (ed.), *New Directions in Semantics*, Academic Press, London, 1987, pp. 1–54.
- Higginbotham, James: 1988, 'Is Semantics Necessary?', *Proceedings of the Aristotelian Society* **88**, 219–241.
- Higginbotham, James: 1989, 'Knowledge of Reference', in George (1989a), pp. 153–174.
- Higginbotham, James: 1989, 'Elucidations of Meaning', *Linguistics and Philosophy* **12**, 465–517.
- Higginbotham, James: 1994, 'Priorities in the Philosophy of Thought', *Proceedings of the Aristotelian Society* **20**(Suppl.), 85–106.
- Higginbotham, James: 1995, 'The Place of Natural Language', in P. Leonardi and M. Santambrogio (eds.), *On Quine*, Cambridge University Press, Cambridge, pp. 113–139.
- Higginbotham, James: 1998, 'Conceptual Competence', in Villanueva 1998, pp. 149–162.
- Higginbotham, James: 1998b, 'On Knowing One's Own Language', in Wright et al. (1998), pp. 429–441.
- Higginbotham, James: 2000, 'On Events in Linguistic Semantics', in James Higginbotham, Fabio Pianesi, and Achille C. Varzi (eds.), *Speaking of Events*, Oxford University Press, Oxford.
- Johnson, Kent: 2002, 'Knowledge of Syntax and Knowledge of Semantics', paper presented at the APA Eastern Division Meeting.

- Johnson, Kent and Ernie Lepore: 2000, 'Knowledge and Semantic Competence', in I. Niiniluoto, M. Sintonen and J. Wolenski (eds.), *The Handbook of Epistemology*, Kluwer Academic Publishers, Dordrecht.
- Kirsh, David: 1990, 'When is Information Explicitly Represented?', in P. Hanson (ed.), *Information, Language, and Cognition*, University of British Columbia Press, Vancouver, pp. 340–365.
- Larson, Richard: 1988, 'On the Double-Object Construction', *Linguistic Inquiry* **19**, 335–392.
- Larson, Richard and Gabriel Segal: 1995, *Knowledge of Meaning*, MIT Press, Cambridge, MA.
- Lepore, Ernest: 1996, 'Conditions on Understanding Language', *Proceedings of the Aristotelian Society*, pp. 41–60.
- Leslie, Alan M.: 1982, 'The Perception of Causality in Infants', *Perception* **11**, 173–186.
- Leslie, Alan: 1995, 'A Theory of Agency', in Dan Sperber, David Premack, and Ann James Premack (eds.), *Causal Cognition*, Clarendon Press, Oxford, pp. 121–141.
- Leslie, Alan and Stephanie Keeble: 1987, 'Do Six-Month-Old Infants Perceive Causality?', *Cognition* **25**, 265–288.
- Levin, Beth: 1993, *English Verb Classes and Alternations*, University of Chicago Press, Chicago.
- Ludlow, Peter: 1999, *Semantics, Tense and Time*, MIT Press, Cambridge, MA.
- Lycan, William: 1986, 'Tacit Belief', in R. J. Bogdan (ed.), *Belief*, Clarendon Press, Oxford, pp. 61–82.
- Marantz, Alec: 1995, 'The Minimalist Program', in Webelhuth (1995), pp. 349–382.
- Perlmutter, David: 1978, 'Impersonal Passives and the Unaccusative Hypothesis', in *Proceedings of the Fourth Annual Berkeley Linguistics Society*, Berkeley Linguistics Society, Berkeley, pp. 157–189.
- Peacocke, Christopher: 1998, 'Implicit Conceptions, Understanding and Rationality', in Enrique Villanueva (1998), pp. 43–88.
- Pettit, Dean: 2002, 'Why Knowledge is Unnecessary for Understanding Language', *Mind* **111**(433), 519–550.
- Pietroski, Paul M.: 1998, 'Actions, Adjuncts, and Agency', *Mind* **107**(425), 73–111.
- Pinker, Steven: 1979, 'Formal Models of Language Learning', *Cognition* **7**, 217–283.
- Pinker, Steven: 1989, *Learnability and Cognition*, MIT Press, Cambridge, MA.
- Pinker, Steven: 1994, *The Language Instinct*, Harper, New York.
- Pinker, Steven: 1999, *Words and Rules*, Basic Books, New York.
- Powers, Lawrence H.: 1978, 'Knowledge by Deduction', *The Philosophical Review* **LXXXVII**(3), 337–371.
- Quine, W. V.: 1960, *Word and Object*, MIT Press, Cambridge, MA.
- Quine, W. V.: 1972, 'Methodological Reflections on Current Linguistic Theory', in D. Davidson and G. Harman (eds.), *Semantics of Natural Language*, D. Reidel, Dordrecht, pp. 442–454.
- Rorty, Richard (ed.): 1967, *The Linguistic Turn*, University of Chicago Press, Chicago.
- Schiffer, Stephen: 1987, *Remnants of Meaning*, MIT Press, Cambridge, MA.
- Sellars, Wilfred: 1956, 'Empiricism and the Philosophy of Mind', in H. Feigl and M. Scriven (eds.), *Minnesota Studies in the Philosophy of Science*, Vol. 2, University of Minnesota Press, Minneapolis.
- Sellars, William: 1969, 'Language as Thought and Communication', *Philosophy and Phenomenological Research* **29**, 000–000.

- Smith, Barry: 1998, 'On Knowing One's Own Language', in Wright et al. (1998), pp. 391–428.
- Stich, Stephen P.: 1971, 'What Every Speaker Knows', *Philosophical Review* **80**, 476–496.
- Stich, Stephen P.: 1978, 'Beliefs and Subdoxastic States', *Philosophy of Science* **45**, 499–518.
- Tanenhaus, Michael, Julie Boland, Gail Mauner, and Greg Carlson: 1993, 'More on Combinatory Lexical Information: Thematic Structure in Parsing and Interpretation', in G. Altmann and R. Shillcock (eds.), *Cognitive Models of Speech Processing*, Lawrence Erlbaum, Hillsdale, NJ, pp. 297–319.
- Tenny, Carol L.: 1994, *Aspectual Roles and the Syntax-Semantics Interface*, Kluwer Academic Publishers, Dordrecht.
- Ullman, Michael T., Suzanne Corkin, Marie Coppola, Gregory Hickok, John H. Growdon, Walter J. Koroshetz, and Steven Pinker: 1997, 'A Neural Dissociation within Language: Evidence that the Mental Dictionary is Part of Declarative Memory, and that Grammatical Rules are Processed by the Procedural System', *Journal of Cognitive Neuroscience* **9**(2), 266–276.
- Urmson, J. O.: 1965, 'J. L. Austin', Reprinted in Richard M. Rorty (1967), pp. 232–238.
- Weberhuth, Gert (ed.): 1995, *Government and Binding Theory and the Minimalist Program*, Blackwell, Oxford.
- Wright, Crispin, Barry Smith, and Cynthia MacDonald (eds.): 1998, *Knowing Our Own Minds*, Clarendon Press, Oxford.
- Vendler, Zeno: 1967, *Linguistics in Philosophy*, Cornell University Press, Ithaca.
- Villanueva, Enrique (ed.): 1998, *Concepts*, Ridgeview Publishing Co., Atascadero, CA.
- Wright, Crispin: 1986a, 'Theories of Meaning and Speaker's Knowledge', in *Realism, Meaning and Truth*, 2nd edn., Blackwell, Oxford (1993), pp. 204–238.

University of California Irvine
Department of Logic and Philosophy of Science
3151 Social Sc. Plaza A
Irvine, CA 92697-5100
USA
E-mail: johnsonk@uci.edu

