

Tacit and accessible understanding of language

Kent Johnson

Received: 19 July 2004 / Accepted: 22 March 2006
© Springer Science+Business Media B.V. 2007

Abstract The empirical nature of our understanding of language is explored. I first show that there are several important and different distinctions between tacit and accessible awareness. I then present empirical evidence concerning our understanding of language. The data suggests that our awareness of sentence-meanings is sometimes merely tacit according to one of these distinctions, but is accessible according to another. I present and defend an interpretation of this mixed view. The present project is shown to impact on several diverse areas, including inferential role semantics and holism, the nature of learning, and the role of linguistics in the law.

Keywords Cognitive science · Psychology · Linguistics · Tacit knowledge · Philosophy of language

Introduction

A central issue in the philosophy of language concerns speakers' epistemic relation(s) to their language. Sometimes the issue appears as an abstract question like "What does our understanding of language consist in?" (Davidson, 1967a, 1970b, 1974; Dummett, 1975, 1978; Peacocke, 1975; Wittgenstein, 1953; Wright, 1986) Other times the question is more concrete: 'What are the details of the workings of a language?' (e.g., Davidson, 1967b, c, 1970a; Higginbotham, 1985, 1989; Parsons, 1990; Pietroski, 1998; this is also the standard project of linguistics, e.g. Larson & Segal, 1995) In this paper, I will steer

I am indebted to a number of people for their useful feedback, especially Peter Ludlow, Paul Pietroski, and two anonymous reviewers. Earlier versions of this paper were presented at an Eastern meeting of the APA, a meeting of the Society for Exact Philosophy at Simon Fraser University, and at a semantics workshop in Ottawa, Canada. I greatly appreciate the comments from those audiences.

K. Johnson (✉)
University of California, Irvine, CA, USA
e-mail: johnsonk@uci.edu

a middle course between these two approaches and ask “What is our cognitive/epistemic relation to the structure of language?” This question is usually addressed by inquiring into the nature of our grasp of a grammar of our language (George, 1989; Nagel, 1969). (As I use the term, a *grammar* is a formal system that characterizes which objects are sentences of the language, along with their pronunciation and interpretation.) Although the question of our relation to grammar is interesting and important, I will focus on a different epistemic phenomenon. In particular, I will inquire into our cognitive grasp of individual sentences: Do we actually understand what our sentences mean?¹ I approach this question by developing an empirically motivated framework that sharpens the issue into a class of empirical questions. I then present some empirical data I have collected that directly address some of these questions. Although the data were mixed, attention to the relevant underlying conceptual principles exposes a consistent, multi-faceted view of our grasp of the meaning and syntax of our sentences. I argue that not only is there an important sense in which we are sometimes only tacitly aware of some elements of meaning, but also that there is another important sense in which we are not. In contrast, our awareness of some basic syntactic features of language appears to be robustly tacit. We’ll also see how these results impact on a number of philosophically important ideas. For instance, I’ll argue that the data create a real problem for some common forms of semantic theorizing, viz., inferential role semantics and semantic holism. We’ll also see how the data interact with a number of other issues, from psychology and linguistics to certain legal issues.

An advantage of my approach is that it allows some philosophically interesting issues to be addressed in part by empirical methods. This result, I maintain, is beneficial both to philosophy and to the relevant empirical sciences. It is beneficial to philosophy because it helps us to hone in on the *exact* content of some philosophical theses, and it exposes novel methods for exploring these theses. It is also useful to the empirically based cognitive sciences, because it helps to raise a number of novel questions for those fields to explore, and it can help increase the precision of some issues and concepts of these fields. Finally, as I discuss below, focusing the general question of our knowledge of language around our grasp of individual sentences is in line with current linguistic theory, which posits very few grammatical “rules” or processes, and accounts for the (apparent) complexity of language by adding much structure to expressions (e.g., Chomsky, 1995).

The paper is organized as follows. In Sect. 1, I discuss the distinction between tacit versus accessible awareness of parts of our environment. I show that there is an entire spectrum of possible (and useful) distinctions here. In Sect. 2, I describe the particular linguistic features that I used to explore our awareness of syntactic and semantic features. In Sect. 3, I describe the experiments I ran and report the results. In Sect. 4, I interpret these results. We see there that which tacit/accessible distinction is used is highly relevant to determining whether or not our awareness of meaning is tacit. In Sect. 5, I discuss some ways this project can be refined and developed.

¹ From an abstract point of view, the fact that the grammar contains a good deal of structure is logically distinct from the fact that our sentences have so much structure. It would be easy, e.g., to construct either an intricate grammatical system that yielded sentential objects containing little or no structure, or a very simple grammar that yields massively structured sentences. Moreover, the structures of a grammar and the sentences it yields are also epistemically independent—the degree and manner of one’s epistemic grasp of a grammar does not necessarily correlate with the degree and manner of one’s epistemic grasp of the structure of the sentences of the language.

I also discuss how it interacts with other philosophically interesting issues. I conclude in Sect. 6.

Before beginning, several important caveats and terminological notes are in order. First and foremost, I must emphasize that this essay borrows from the methods (and to some extent the concerns) of linguistics and psychology, as well as philosophy. Fortunately, I will be able to introduce and motivate these methods as the project is described. At the same time, I do not cover every detail of the relevant issues from all the various perspectives: one paper can scarcely do justice to the topic of linguistic awareness from the perspective of one field, much less three of them. Rather, my aim here is to describe a general project and to raise a number of interesting multidisciplinary questions.

A few minor notes: I will speak of tacit and accessible “awareness”, instead of using the more familiar labels of tacit and accessible knowledge or belief. My reasons for doing so will become clearer later on. For now though, suffice it to say that the relevant epistemic notion I use is more akin to perception than to propositional knowledge. Also, when speaking of semantics and semantic features, I use these notions in an “internalist” sense. That is, I focus only on the mental representations (e.g., “concepts”) associated with sentences that enable us to pick out entities in the world. I do not address any issues of how our words hook up to these entities. (I do this primarily to avoid becoming embroiled in metaphysical debates about the semantic elements I employ. Whatever their metaphysical/referential status, the elements I discuss play an important role in our cognitive lives, particularly as concerns our linguistic abilities, and it is this role that I aim to explore.)

1 Tacit and accessible awareness

What is the difference between being tacitly aware of something, and having a more robust awareness of it? When speaking of tacit awareness, we wish to distinguish those features (of our environment, e.g.) which we “can”—in some appropriate sense of this term—become explicitly aware of so that it “can” figure into our overt rational thoughts and deliberations. We can codify the idea behind this distinction with the following criterion of accessibility:

- (1) A person is accessibly aware of some feature if and only if she can demonstrate explicit awareness of the feature after she has been given an appropriate sort of priming. (By stipulation, a person is tacitly aware of a feature iff her awareness is not accessible to her.)

((1) is not intended as a conceptual analysis of any sort. It is used only to gesture at the relevant distinction, with fine-tuning to be expected as an empirical matter of course.) “Priming” here is a blanket term for the background stimuli or information that helps to direct the person’s attention. This background information also includes how a question is asked—e.g., *What do you think the answer is?* is a different form of priming than *Is the answer 2π?* The criterion in (1) raises two questions: What is to count as an appropriate sort of priming, and what is it to be able to demonstrate explicit awareness of a feature? I take these questions in turn.

What sort of priming is appropriate for deciding whether a subject has tacit or accessible awareness of a feature of her environment? Crucially, there is no single answer to this question, because a lot depends on the kind of project at hand. At one

extreme end, no priming whatsoever should be given to a subject. For instance, if you are testing a hypothesis that the order in which a person lists various objects (tools, pieces of furniture, etc.) varies according to the person's gender, you would not want to give subjects much background information at all. Rather, you might simply show them a visual array and ask them to list the objects they see there. Similarly, in a well-known series of experiments, subjects were asked to rate a number of photographs of female faces in terms of their attractiveness. Unbeknownst to the subjects, one photo appeared twice in the stack, except that the pupils in the photo had been dilated, indicating excitement (Hess, 1965, 1975; Stich, 1978). In order to test whether subjects would be aware of this difference, it was important not to give them any priming to attend to pupils in the photos. (Incidentally, subjects tended not to list pupil size as a reason for finding the doctored photo more attractive.)

On the other hand, it is sometimes important to give subjects quite a bit of priming. If, for instance, you are exploring subjects' abilities to manipulate sounds of their language, you may need to help them see the different phonemes. Ask a typical speaker how many sounds of English are represented by the letter *t*, and her response may well be that there is just one sound, *tuh*. But a speaker can quickly be led to see that *t* represents many sounds in English: the *t* in *Batman* is typically not pronounced like *tuh*, but instead represents a short burst of silence, like the burst between syllables in *uh-oh*. (Phonologists call this sound a *glottal stop*.) Also, the *ts* in *tip* and *stab* are different, insofar as the *t* in *tip* is aspirated, requiring a short puff of air for its proper pronunciation, but the *t* in *stab* is unaspirated, omitting any such puff. Once these distinctions have been brought to speakers' attention, they may be better prepared to perform the test at hand. The literature on blindsight provides another example of this sort. Blindsight is a disorder which appears to blind a portion of the normal visual field. However, it's known that persons with blindsight receive information from their blindsighted visual region, although they are unable to make use of this information in overt conscious reasoning and deliberations (e.g., Marcel, 1989; Schacter, 1989; cf. Block, 1995 for philosophical discussion). For example, in some experiments, subjects were told that there was an image of a rectangle in their blindsighted region, and they had to guess whether the rectangle was horizontally or vertically oriented. Surprisingly, subjects did significantly better than chance in these tests. Notice, though, that in order to perform the tests, subjects were given the maximum possible amount of priming: they were told everything about what was in their blindsighted region except for the orientation of the rectangle. (It's this large amount of priming that reduces the task to the familiar two-alternative forced-choice paradigm.) Without all this priming, the subjects would not produce a significant number of correct guesses (Marcel, 1989, 1993; Vision, 1998).

These examples show that no single type of priming determines whether or not we are only tacitly aware of something. Different projects employ different amounts of priming, depending on the type of cognitive activity or process they are designed to explore. So two different projects may legitimately use distinct tacit/accessible distinctions if the projects differ about what counts as appropriate priming. It's worth noting that other authors have found similar distinctions to be important as well. For instance, Higginbotham (1998) explores various characterizations of what he calls the conception of a concept, where these characterizations turn on (roughly speaking) what it takes to deliver information about the concept. Later on, I will have more to say about the complexity of this dimension of variation in the tacit/accessible distinction. For now, though, I turn to the second question about (1).

What counts as demonstrating explicit awareness of a feature? Here too, different criteria may be appropriate in different circumstances. In the present case of language understanding, a natural thought is that a subject demonstrates explicit awareness of a semantic feature by accepting inferences that crucially rely on the presence of the feature, and by denying inferences when the feature is absent (this common idea has roots going back at least to Wittgenstein, 1953). While this proposal will work for semantic features, it has no obvious correlate in the realm of syntax. Of course, one might accept or deny an inference based on the *meaning* of a syntactic feature, but the present project is to see if meaning and syntax can be prized apart. So instead of looking to inferential patterns, I suggest that subjects can demonstrate explicit awareness of a linguistic feature by identifying the feature in linguistic environments that are unlike any environments that are used to prime the subject to attend to the target feature. (Notice that a linguistic “environment” will in fact be inside the speaker’s head; it is the result of the speaker having processed a given sentence.) This proposal for demonstrating explicit awareness applies in precisely the same form when we are exploring our awareness of syntactic features. It also retains the spirit of the previous proposal, because if a subject can identify a semantic feature in an environment, she should also be able to accept or deny the corresponding inferences.

Much more could be said about the differences between tacit and accessible awareness. However, the two points just made—that the notion of legitimate priming can vary from project to project, and that accessible awareness can be demonstrated by identifying features in novel environments—will suffice for now. In the next section, I introduce the semantic and syntactic features that I compared in terms of their respective degrees of accessibility.

2 Target semantic and syntactic features: Direct causation and traces

There are too many semantic and syntactic features of language to explore our awareness of each one. So instead, I worked with two representative features: the semantic feature of direct causation, and the syntactic feature of a trace. In this section I describe these features in turn. Importantly, I adopt without defense a number of linguistic assumptions about certain constructions. These assumptions are quite commonly made in the literature but others may nonetheless question them. But if, as I also assume, linguistic theories should be taken seriously as characterizations of a particular psychological ability, we should be prepared to take such assumptions seriously as we find them in the literature. To do otherwise would be to expose linguistic theory to the risk of becoming scientifically insular and isolated. Thus, the present project may be seen as the beginnings of a reconciliation of what David Plaut has described as “the considerable divergence of goals and methods between linguistic versus psychological approaches to the study of language” (Plaut, 2003, 146).

The first feature I’ll discuss is direct causation, which is found in various grammatical constructions. Most famously, it appears in “causative” verbs like *break*, *melt*, *sharpen*, etc. When these verbs occur in their transitive form, they have roughly the meaning that the subject of the verb causes the object of the verb to break, melt, become sharpened, etc.

- (2) Jane caused the vase to break.
- (3) Jane broke the vase.

Sentences (2) and (3) are quite similar in meaning. However, they don't mean exactly the same things. (3) entails not only that Jane caused the vase to break, but also that Jane was "directly" responsible for the vase's breaking. We can see this in the following scenario:

Jane runs an antique store, and she just bought an expensive vase. Jane also has an assistant Dave, who is a trouble-maker. Dave didn't know anything about the vase until Jane mentioned it to him. When Jane told Dave that he should be careful with the new vase because it was so valuable, he decided to drop the vase on the cement floor. When the vase hit the cement, it shattered into a thousand pieces.

In this scenario, did Jane cause the vase to break? In one sense, she did, since Jane did do something—she warned Dave about the vase—and this made Dave decide to drop the vase. Had she not said anything to Dave, he would not have touched the vase. So in the above scenario, sentence (2) is true (cf. Lewis, 1973). But (3) is false: Jane didn't break the vase, Dave did. In order for (3) to be true, the subject of the verb (Jane) must directly cause the end result (the breaking of the vase). The sort of causation relevant to the truth of sentences like (3) is called "direct causation". Importantly, the relevant notion of "directness" is a psychological-cum-linguistic one. Our language is structured so as to recognize a certain individuating taxonomy of events. According to this taxonomy, Dave's pushing the vase to the floor was the event that directly caused the vase to break. However, Jane's warning Dave only indirectly caused the vase to break, because the connection between these two events was mediated by at least one more event (cf. Bittner 1999 for detailed linguistic discussion; cf. Pietroski, 1998, 2000 for discussion of the relation of direct causation to the philosophical notion of agency). That this notion of unmediated causation might not jibe with a metaphysical theory is, strictly speaking, irrelevant here.

I assume that direct causation is a genuine semantic feature of many sentences. (As I use the term, a "semantic feature" is a determinate semantic constituent that is posited by our best linguistic theories as part of the total meaning of some sentences. So although direct causation is probably a semantic feature, the property of entailing maleness is probably not a semantic feature of the word *bachelor*.) This claim has been widely endorsed by the linguistic community. Direct causation has played an explanatory role in linguistic studies of such languages as Berber, Chukchee, Icelandic, Turkish, Russian, Tagalog, Malagasy, St'át'imcets, Papago and 'many hundreds of other languages' where it occurs as an overt morpheme attached to a verbal root (Baker, 1988; Bittner, 1999; Hale & Keyser, 1993, p. 102; cf. e.g., Comrie, 1985; Tenny, 1994). For instance, in Japanese, direct causation is expressed by the affix *ase*:

- (4) a. 3-tu-no iyana koto-ga Hanako-o kanasim-ase-ta.
 b. 3-cl-GEN terrible thing-NOM Hanako-ACC sad- CAUSE-Past
 c. 'Three terrible things saddened Hanako.'

In short, it is almost beyond question that direct causation is a part of the grammar of human languages.² (To remind: as I am using the term, a grammar is responsible for more than just the syntax of natural languages; it is also responsible for the semantic (and phonological, etc.) properties of language. So my claim that direct causation

² Fodor (e.g. 1998, ch. 3) is perhaps the only person who denies this last claim. I will not discuss his position here; cf. Johnson, 2004a, b for more detailed treatments of the relevant positions.

is a grammatical feature of human languages does not imply that it must always be realized morpho-syntactically, although it may be). Direct causation is expressed not only in causative verbs, but also in the resultative construction, as in (5a):

- (5) a. Mary shot John dead.
b. Mary's shooting John caused him to become dead.

In its normal interpretation, (5a) means something like (5b), but the kind of causation required must be direct when (5a) is true. Suppose Mary shoots and hits John on the tip of his pinkie finger. Although the wound is superficial, the sight of his own blood terrifies John so much that he has a heart attack and dies (cf. Bittner, 1999). In such a case, although (5b) is true, (5a) is clearly false (for further discussion, cf. e.g., Baker, 2003; Bittner, 1999; Levin & Rappaport Hovav, 1995).

Three reasons conspire to make direct causation a nice feature to work with. First, we just saw two different sorts of constructions—causative verbs and resultatives—which contain it. So subjects can be primed to attend to direct causation using one sort of construction, and their awareness of this feature can be tested using the other type of construction. Second, in both these constructions the causative morpheme is not expressed either phonologically, in writing, or in some overt syntactic way. So if subjects can identify direct causation in novel environments, that will be evidence that they are really spotting a semantic feature (and not just some phonological or syntactic reflex of the feature). Third, as (6) shows, some sentences are superficially similar to sentences containing either causatives or resultatives:

- (6) a. Jane broke the vase.
b. Jane likes the vase.
c. Mary shot John dead.
d. Mary ate the meat raw.

There is no direct causation expressed by the verb *likes* in (6b), and (6d) only expresses that Mary ate the meat while the meat was in the state of being raw. Thus, if subjects can detect direct causation, they should be able to distinguish cases like (6a) and (6c) from cases like (6b) and (6d).

I contrasted the semantic feature of direct causation with the syntactic feature of a trace (or *wh*-trace in particular).³ To see what a trace is, consider the following two sentences:

- (7) a. Melanie kicked John.
b. Who did Melanie kick?

In (7a), *John* is the object of *kick*; however, in (7b), the object is *who*. Even though *who* is in the front of the sentence, it still has an important structural relationship to the object position right after the verb. For example, you cannot simply add an object to (7b): **Who did Melanie kick John?* is ungrammatical. We can think of the word *who* as

³ As a technical caveat, I here assume a very basic form of Government and Binding Theory syntax (Chomsky, 1981). I choose it because it is the most well-known theory of syntax, and many other styles of syntactic theory make corresponding assumptions which, for present purposes, are equivalent to my use of *wh*-traces. (For example, many contemporary Minimalist theories are effectively equivalent to the one I use, since the result of an operation like Copy (possibly with Delete) produce constituents we may treat like traces.) Moreover, the general point illustrated with traces should be replicable (*mutatis mutandis*) with many other syntactic features.

“moving” from the object position of the verb *kick* and eventually landing at the front of the sentence. Each time *who* makes a move toward the front of the sentence, we can think of it as leaving a trace of itself, which occupies the position it leaves. So when *who* leaves the object position of *kick*, a trace of it is still present there. The presence of this trace means that the position cannot be filled in with another word like *John*. In more complex sentences, we get more traces. When a *wh*-word is extracted from an embedded clause, there is a trace present at the front of the embedded clause:

(8) Who do you think *t* Melanie likes *t*?

(9) gives a bit of evidence for this “intermediate” trace. Although embedded questions can have a *wh*-word appearing at the front of an embedded clause, as in (9a), when another *wh*-word moves from the embedded clause to the front of the entire sentence (9b), the front of the embedded clause can no longer host a *wh*-word, as (9c) shows.

- (9) a. Stacy knows when₁ Melanie kissed John *t*₁.
 b. Who₂ does Stacy know *t*₂ Melanie kissed *t*₂ last week?
 c. *Who₂ does Stacy know when₁/*t*₂ Melanie kissed *t*₂ *t*₁?

The discussion of (8)–(9) is not intended as a detailed characterization or defense of traces in human languages. Such a task is far beyond the scope of this paper. Instead, I only mean to introduce the notion of a trace, and I will simply follow a very standard view and assume that traces (or something like them) really are present in the positions indicated (Chomsky, 1981, 2000, 2001; Culicover, 1997; Haegemann, 1994). (For those who are not friends of traces, many other syntactic notions could be used, with much the same outcome; see below for further discussion of this matter; cf. also the previous footnote.)

Traces are useful syntactic entities for our present purposes, because in many ways they are precisely analogous to the semantic entity direct causation. Traces appear in at least two different environments—object position, and the front of embedded clauses. Moreover, they are phonologically and orthographically unrealized. Also, the intermediate trace at the front of an embedded clause has no obvious semantic role to play, so a trace in this position cannot be detected by attending to the meaning of the sentence and noticing a logical gap that needs to be filled in by the *wh*-phrase. Together, direct causation and traces provide just the sort of linguistic materials needed to examine our awareness of syntactic and semantic properties. I will have more to say about my choice of these two elements of linguistic theory later, but for now I turn to the experiments.

3 Empirical evidence

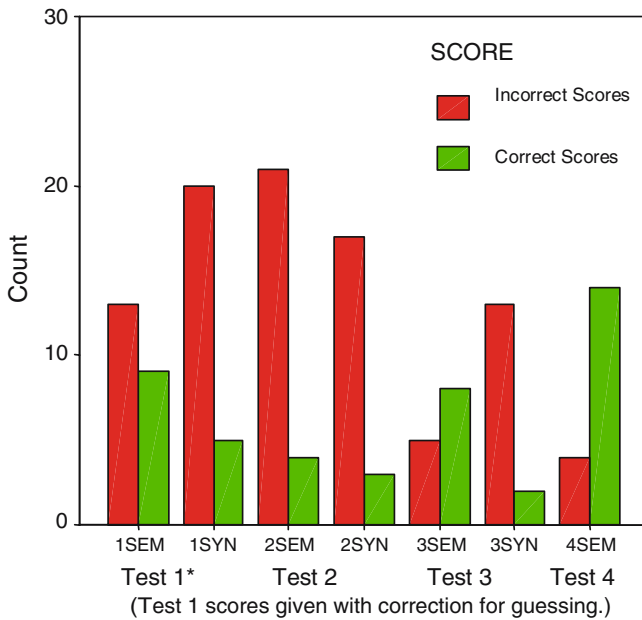
I tested our awareness of semantic and syntactic features by giving students at UC Irvine and Cal State Long Beach a short questionnaire. The questionnaire began with a few background questions to identify those students who were native speakers of English and who had no background in linguistics. Only these subjects were used in the study. Some subjects read a characterization of direct causation (the “semantic condition”), while others read a characterization of traces (the “syntactic condition”). The semantic condition described examples like (2)–(3) above, contrasting causative verbs with sentences like *Jane likes the vase*. (Most of the paragraph containing

(2)–(3) above is taken verbatim from the actual questionnaire.) The syntactic condition discussed traces with examples like (7), and they were contrasted with sentences like *Bill is happy* and *Jane thinks that Bob is silly*. These characterizations were each about 380 words in length (the length of the characterization of traces was always slightly longer than the characterization of direct causation), and they were similar in content to the discussions in Sect. 2. The description of direct causation described the semantic feature purely semantically; no reference to the structure of the sentence was made. In contrast, the description of traces was purely syntactic, with no mention made of the interpretation of the sentences.⁴ After reading the characterization, subjects then turned to a list of five sentences, and they were asked to indicate where, if anywhere, the target feature—direct causation or traces—occurs in these sentences. Subjects were instructed to adopt the most natural, normal readings of the sentences in question, under which they might be plausibly used in everyday conversation. The five sentences were: two screening sentences, two target sentences, and one filler. They were taken randomly from a set of 20 screening sentences (10 of each kind, cf. below), 20 target sentences (10 of each kind) and 10 filler sentences. The sentences were presented in random order, except that the two screening sentences always appeared first on the list. The two screening sentences consisted of two sentences that presented the target feature’s presence or absence in the same way that the examples illustrated the feature in the priming (e.g., *Joan flattened the box*, *Peter hates Mary*, when subjects were primed using (2)–(3)). Only subjects who correctly identified the presence or absence of the feature in question were scored. There were two target sentences, a “target-yes” sentence and a “target-no” sentence. In the semantic condition, the target-yes sentence expressed direct causation in a resultative construction like *Mary shot John dead*. The target-no sentence had a similar form, but it was not a resultative, e.g., *Mary ate the meat raw*. Under its most normal reading, this last sentence means that Mary ate the meat while the meat was in the state of being raw. (Although the sentence does have a nonsensical causative reading, this is clearly not a normal use. All sentences used had only one relevant reading that is normal, and subjects were explicitly instructed to adopt the most natural reading of the sentence.) In the syntactic condition, the target-yes sentence contained a trace in the front of an embedded clause, as in *Who do you think Melanie likes?* The target-no sentence was an unrelated filler sentence, such as *Karen knew the answer*. Thus, in each condition, the subject had to analyze the sentence (semantically or syntactically) and determine whether the appropriate element was present (in the target location).

Subjects’ responses were scored based on whether they correctly indicated the presence of the linguistic feature in the target-yes sentence, and whether they correctly indicated its absence in the target-no sentence. (They indicated this by marking either the *Yes* or the *No* beside each sentence.) Scores were then grouped into two categories: those that had correct responses for both target questions (a “correct” score), and those that had at least one error (an “incorrect” score). Thus, in all cases, the main score was based on the responses to a pair of binary forced-choice alternatives. (Hence chance of getting a correct score by random guesses was .25.) So in each test, each subject fell into one of four categories, depending on whether they were in the semantic or syntactic condition and whether or not their score was correct. The

⁴ To emphasize the syntactic nature of traces, subjects were also told that subject *wh*-words (as in *Who kissed Ted?*) also leave a trace in the subject position.

results of the four main tests are summarized in the figure below. (A few ancillary experiments are not presented here, but are presented in the text.)



In Test 1, subjects' attention was directed towards the relevant portions of the sentences by underlining. Subjects saw sentences like *Mary shot John dead*, (or *Who do you think Melanie likes?*). Subjects were asked to indicate whether or not direct causation (or a trace) occurs in the underlined portion of the sentence. (Each filler sentence had some arbitrary part underlined.) 54.5% of the subjects had correct scores in the semantic condition, and 40% of the subjects had correct scores in the syntactic condition, too. It seemed excessive to suppose that 40% of the subjects would know that a trace is left at the top of an embedded clause, and an examination of the wording of the instructions suggested that subjects simply made random guesses when they did not know whether or not a trace was present. A correction for guessing reduced the percentage of correct scores in the syntactic condition to 20%. This score was in line with the other scores for the syntactic condition, and these other scores remained constant throughout the various experiments. Thus, I assumed that the first scores reflected guessing. But then a correction for guessing was required in the semantic condition, too, which reduced that percentage of correct scores to 39.4%, which is not significantly different from the syntactic condition ($N = 47$, $\chi^2 = 1.37$ (1 df), $p > .11$, with correction for guessing; without correction for guessing, the results were still not significant: $\chi^2 = .994$ (1 df), $p = .319$).⁵ In the remaining experiments, the instructions were changed, and subjects were told to indicate that the linguistic feature was present only if they thought it was, and if they had no opinion, they were to indicate that it was not present, rather than just guessing randomly. Thus, correction for guessing was used only in this first experiment.

In Test 2, subjects were asked to provide a written gloss of each sentence before indicating whether the target linguistic feature was present or not. (In the two conditions,

⁵ As is standard, N indicates the number of subjects in the experiment.

subjects were asked to pretend that they were explaining the sentence to someone who does not understand “what it means” (semantic condition) or “how it is put together” (syntactic condition.) Subjects were instructed to pay special attention to whether the target feature occurred anywhere in the sentence. The relevant portions of the sentence were not underlined; instead, subjects were simply asked to mark where, if anywhere, in the sentence the linguistic feature occurred.⁶ Here subjects fared about the same in the syntactic condition (15% correct), but their performance in the semantic condition was significantly worse than in the first experiment (16% correct, $N = 47$, $\chi^2 = 7.74$ (1 df), $p < .01$). Ultimately, though, the difference between the syntactic and semantic conditions in Test 2 was not significant ($N = 45$, Fisher’s Exact Test (1 df), $p > .6$).

In Test 3, subjects were asked to gloss each sentence, and the relevant portion of each target sentence was underlined. Furthermore, each target sentence was presented as part of a contrasting environment. E.g., subjects received sentences like “*Jacob only wiped his table a little bit and it was still dirty, but Mary wiped her table clean*”, and “*Who does Peter like, and who do you think Melanie likes?*”. (Subjects were only asked to gloss the unitalicized portion of the sentence.) Collectively, these techniques made a difference. The percentage of correct scores in the syntactic condition was 13.3% but 61.5% of the scores in the semantic condition were correct. The distinction between the syntactic and semantic conditions was significant ($N = 28$, Fisher’s Exact Test (1 df), $p = .011$).

In Tests 2 and 3, many subjects who responded incorrectly to the target-yes sentence (by marking the *No* alongside the sentence) had glossed it with temporal vocabulary. E.g., they often explained *Mary wiped the table clean* as “Mary wiped the table until it was/became clean”. This tendency raised the question of whether an added stress on the notion of “becoming” would increase the number of correct scores. So I tried the semantic condition of the third experiment again (as Test 4) with an increased usage of words like “become”. For instance, direct causation was characterized as when one thing directly causes another else to become something, where the target example *Jane broke the vase* was glossed as meaning that Jane caused the vase to become broken. In this semantic condition, the percentage of correct scores was increased to 77.8%, which did not differ significantly from the semantic condition in Test 3 ($N = 30$, Fisher’s Exact Test (1 df), $p > .27$). However, the contrast between the semantic condition of Test 4 and the syntactic condition of Test 3 was very significant ($N = 33$, $\chi^2 = 13.6$ (1 df), $p < .0003$). The fact that the strategy of Test 4 seemed to improve speakers’ abilities to recognize direct causation is especially interesting, since many linguists have posited an element of becoming in the semantic representation of causative constructions (e.g., Baker, 2003; Dowty, 1979; Hale & Keyser, 1987; Parsons, 1990; Pustejovsky, 1991).

The four main tests discussed above were supplemented with several ancillary tests. To get some idea of the advantage provided by the written characterization of direct causation, some subjects were given a test exactly like Test 1, except that no charac-

⁶ There was an extra amount of space left between each word to indicate that a linguistic feature occurred between two words, and subjects were allowed to indicate general areas where the feature occurred, if they could not pinpoint its exact location. These tests were scored very liberally; in particular, an answer was marked as correct in the target-yes question if there was any indication that the feature was present in the given location. Subjects were given an example of how to mark their guesses, and there were no clear cases of difficult judgments as to whether subjects had correctly identified the target feature.

terization of direct causation was included at all. Subjects were only told that direct causation is the sort of meaning expressed in a sentence like *Dave broke the vase*, but which is not expressed in *Dave caused the vase to break*. Compared to Test 1, subjects' correct scores were reduced from 39.4 to 6.7% (after correction for guessing), and this difference was significant (Test 5: $N = 42$, $\chi^2 = 4.52$ (1 df), $p < .05$; as a common measure of effect size here, we have $\phi = .328$). So the priming was evidently useful. Furthermore, alongside of Tests 1 and 3, I also gave some subjects a similarly structured test, except that the priming used a resultative construction, and the target-yes and target-no sentences were causative and non-causative transitive verb construction respectively (cf. (6a–b)). In these variants, subjects' performance was consistently better. In Test 1, switching causatives and resultative improved scores from 39.4% to 50% (Test 6: $N = 46$, $\chi^2 = .382$ (1 df), $p = \text{ns}$). In Test 3, this switch improved scores from 61.5 to 78.6% (Test 7: $N = 25$, Fisher's Exact Test (1 df), $p = .293$).

4 Interpreting the evidence

How should we understand the data presented in Sect. 3? Although the evidence only concerns one semantic feature and one syntactic feature, in this section I will explore what it suggests about our awareness of semantics and syntax more broadly. The evidence suggests three things. First, it suggests that it's hard to find a significant cognitive distinction between our awareness of syntactic features and our awareness of semantic features when these features are not in an expression's overt structure. Second, subjects' awareness of the syntax of their language is robustly tacit. Third, subjects' awareness of meaning is somewhat malleable, and under the right conditions a distinction in awareness between syntax and semantics can be observed. I discuss each of these claims in turn.

4.1 A sense in which we are only tacitly aware of sentence-meanings

Although we finally drove a wedge between our awareness of syntactic versus semantic features, the initial experiments were unsuccessful, despite containing a lot of priming. Test 2 was an extreme case. There, subjects received two sources of priming—a page-long description of direct causation, and the explicit instructions to write a gloss of the meaning of the sentence, paying special attention to whether direct causation occurs anywhere in the sentence. The priming gave subjects ample opportunity to reflect on and identify the various semantic features of the sentences. The fact that their performance was so bad was surprising, and must be taken seriously (Indeed, in Test 2, the score was slightly below chance, although not significantly (binomial test, $p = .214$.) The crucial interpretive question is whether these tests provide evidence that subjects are only tacitly aware of direct causation. But as we saw earlier—cf. (1)—the difference between tacit and accessible awareness can vary, depending on what counts as legitimate priming. Indeed, it's probably best not to think of there being a “tacit/accessible” distinction at all; explicit reference at least to the type of priming involved must be made. Just as there is no distinction between those numbers n such that $n + x$ is odd/even until a value for x is specified, so too there is no tacit/accessible distinction until at least the nature of the priming is specified.

The importance of specifying the priming used emphasizes the need for it to meet various standards of scientific appropriateness. For many purposes, the amount of

priming used in Test 2 would be excessive. For instance, a corresponding amount of priming would surely enable us to demonstrate explicit awareness that *Do geese see God?* is a palindrome, or that two seemingly identical photos of a woman were different because the pupils in one of them had been enlarged. Yet both of these have been offered as exemplar cases of things that we do not normally believe at all, not even tacitly (Crimmins, 1992; Stich, 1978). Thus, these experiments suggest an important sense in which our awareness of the semantic features of our sentences is not always fully accessible to us. This last claim can be put in a rather striking form. To see this, notice that in one very typical sense of “understanding”, one understands a sentence only if one understands its component semantic features and recognizes their presence in the sentence. For example, if you only understand that *Steve baked Lucy a cake* involves Steve, baking, Lucy, and a cake, and you don’t see that Steve baked the cake, and he did it for Lucy, then in one sense, you do not understand the sentence. Similarly, if you know that this sentence means that Steve baked Lucy a cake, but you think that *Steve* expresses the relation of baking, *a* refers to Lucy, *bake* denotes cakes, *Lucy* denotes the past tense, *-ed* refers to Steve, and *cake* is an indefinite determiner, then you don’t understand the sentence according to this common notion. According to this sense of understanding, then, *we do not always understand the meanings of our own sentences*. That is, we don’t understand the meanings of our sentences in the same way that we don’t understand their syntax (cf. below) or the means by which we make the other sorts of judgments and categorizations mentioned above. Moreover, this claim should be understood internalistically: we are not always aware of the senses that *our very own linguistic abilities* assign to our sentences. In other words, the situation regarding our epistemic grasp of language is worse than the Kripke/Putnam/Burge-style cases suggest. Those cases purport to show that we may not always understand what our sentences mean in the sense that we may not know which things, properties, etc. our expressions pick out. The present work shows that, even when using very generous forms of priming, we may not have access to the internal “senses” or “modes of presentation” (to borrow some neo-Fregean terminology) that are responsible for, and perhaps constitutive of, the corresponding expression’s denoting what it in fact denotes.

Notice, incidentally, that the present issue appears to be quite distinct from issues of processing difficulty. It’s well-known that people have a hard time understanding sentences like *The man the dog the wolf scared bit died* or *No eye injury is too trivial to ignore*. But the trouble with these sentences appears to center around difficulties in putting the sentence together. In the examples we’ve used, there doesn’t appear to be any problem processing the target sentences; rather, the problem seems to turn on observing the presence of a semantic constituent. In other work (Johnson, 2004a), I defended, on largely intuitive grounds, the standard view that there’s a straightforward cognitive distinction between syntactic and semantic properties. While I believe that the claims from that paper still hold (cf. below), the present work helps to expose just how nontrivial such an empirical claim is. Indeed, the cognitive distinction between syntax and semantics is much less dramatic than it initially appears to be.

At first glance, it might seem odd to say that we sometimes don’t understand what our own sentences mean. Although it’s plausible that we only tacitly know a grammar of our language, surely we at least know what our sentences mean? After all, we are the ones who decide what proposition we wish to express and communicate, and we decide which sentence to use to express it. Part of this concern is surely right. There

is an important sense in which we consciously and explicitly decide “what to say”, and we consciously and explicitly decide (at least in part) on the form of words to express it. No normal speaker of English, e.g., attempts to say that the cat is on the mat but mistakenly produces the sentence *snow is white* out of confusion about what the sentence means. Moreover, we are robustly aware of some aspects of meaning. For instance, it would not take much priming at all to get speakers to correctly indicate those parts of the sentence that refer to males (as in *the boy kissed the girl*). But the present issue is not about whether we are competent speakers—we are—or whether there are many semantic aspects of sentences that we are accessibly aware of—there are. The evidence presented above suggests that there are certain “smallish” bits of meaning which, in certain circumstances at least, we are simply not aware of.⁷ This last claim is not odd; in fact, it’s exactly what we should expect. I will explain this from two different perspectives: that of cognitive science in general, and that of psycholinguistics in particular.

From the general perspective of cognitive science, the claim that we have a less than full awareness of our sentences’ meanings is par for the course. The last thirty or forty years have shown that (virtually) all our complex cognitive processes and abilities—including those that seem to be under our rational control or at least transparently introspectible—are grasped much less completely than we pretheoretically tend to think. For instance, it may seem obvious that if there are no Cartesian demons or evil hypnotists around, our clear and distinct memories are straightforward reflections of our past experience. However, humans are surprisingly susceptible to memory implantation and memory manipulation (Loftus & Hoffman, 1989; Loftus & Ketcham, 1994). It is quite easy, for instance, for a person in normal circumstances to receive just a few false suggestions, and later “remember” events that never occurred, including details that were never part of the original false suggestions. We are also often unaware of elements of a visual scene. For instance, when counting the number of passes in a ball game played in a small hallway, many subjects failed to notice that a person in a gorilla suit stepped into the middle of the visual display, thumped her chest, and walked away (Simons & Chabris, 1999). As a final example, it seems obvious that we understand the stream of our own mental processes, but a large literature on judgment and decision-making suggests otherwise. One early study showed that subjects are often unaware of their own mental processes that enabled them to solve a problem only moments earlier (Maier, 1931). This example is quite typical: many other striking limitations on our awareness of paradigmatically rational processes have also been uncovered (Connolly, Arkes, & Hammond, 2000; Gilovich, Griffin, & Kahneman, 2002; Nisbett & Wilson, 1977). Indeed, some anthropologists and evolutionary psychologists have argued that it can be desirable for our conscious, rational deliberative abilities to lack certain types of information about our actual motivations (e.g., Trivers, 1985). In short, when language is viewed as just another complex rational cognitive process, with interpreted sentences being the outcomes of that process, it’s unsurprising that we have an imperfect awareness of the semantic features of sentences.

Although limited awareness of sentence-meaning isn’t out of step with cognitive science in general, perhaps there is something special about language. Does the

⁷ I put ‘smallish’ in scare quotes because it is hard to say what makes a semantic feature smallish other than that we sometimes don’t observe it. I suspect that being a smallish feature correlates very strongly with being a semantically non-vacuous functional element, but I will not argue for that claim here.

linguistic or psycholinguistic evidence suggest that speakers should always have an accessible awareness of sentence-meaning? It appears not. To see this, recall Sachs's classic demonstration that we are very poor at verbatim memory of sentences (Sachs, 1967). Rather than remembering the particular details of a given sentence, we often simply remember the meaning of the sentence. But our ability to remember sentences verbatim is malleable. When there is practical benefit in remembering the exact phrasing of a sentence (e.g., in cases of jokes or personal insults) our verbatim memory improves significantly (e.g., Murphy & Shapiro, 1994). Converse results hold for semantic content: our memory of various parts of the content of a story can be improved or degraded by adopting the perspective of different participants in the story (Anderson & Pichert, 1978; cf. also Owens, Bower, & Black, 1979). These experiments all point toward a lesson that is very familiar in cognitive science: human attentional resources and working memory are limited, and cognitive systems often respect these limitations by compressing information in various ways. In normal circumstances, for instance, there is no practical difference between *She sent a letter to Tim* and *She sent Tim a letter*, and so we do not ordinarily retain information that distinguishes which of these sentences we heard. (In special circumstances, of course, the difference between these two sentences might be crucial.) I suggest that something similar is going on in the present case. It appears that we often don't notice some smallish semantic features of sentences because they are typically not of much practical importance. For example, it may be that we often don't notice the expression of direct causation in resultative clauses because it is normally unimportant. Ordinarily the vital information to extract from *Mary wiped the table clean* is that Mary wiped the table and that the table was clean afterwards. (We saw some evidence for this last claim when I noted that many subjects glossed such sentences as 'Mary wiped the table until it became clean', thus omitting any causal connection between the wiping and the being clean.) That Mary's wiping caused the table to become clean is typically of little importance, and so the hypothesis predicts that we typically do not store the information about direct causation. This proposal looks efficient. Since we do store the information that Mary wiped the table and that the table was clean afterwards, it should be easy to deduce at a later time that Mary's wiping caused the table to be clean, if we ever needed this specific bit of information.⁸ Moreover, such a proposal would have application beyond just the resultative construction. For instance, even trained linguists occasionally think that *The article angered John* and *John was angry at the article* are synonymous. The information that is usually relevant in either sentence is the same. However, they are not synonyms: only the first sentence is true if John thought the article was beautifully and responsibly written, but became angry at the great injustices that the article described. (Obviously, there will be other relevant factors in our sentence-perception. Semantic features that are phonologically realized are undoubtedly more easily incorporated into the overt mental representation of a sentence. The present hypothesis only suggests that when these other factors have negligible impact on processing, issues of relevance may become important.)

In addition to explaining subjects' poor performance on many of the tests described in Sect. 3, my proposal also makes a number of predictions. One such prediction concerns subjects' improved performance on the tests that reversed the roles of re-

⁸ Of course, this is not to say that the sentence doesn't contain direct causation as part of its actual meaning. Rather, a background assumption of the project is that it does; the point here is that not every feature of a sentence necessarily makes it out of the linguistic regions of the mind and into those regions that constitute conscious awareness.

sultatives and causative constructions. When primed on resultatives, subjects were better at spotting direct causation in causative verbs than vice-versa. This can be explained by the fact that direct causation is typically a highly relevant component of the meaning of a clause based around a causative verb. Although a speaker might omit the direct causation when interpreting a resultative, it seems harder for this to happen with causatives. The typical practical purpose of saying e.g., *Jane broke the vase* is to communicate that Jane did something to the vase, namely break it. But to say that Jane did something in this context is to say that Jane directly caused something to happen. If this is right, then my proposal correctly predicts that subjects will fare better at spotting direct causation in causatives than they do in resultatives. Another prediction comes if we assume that the relevant sentence-processing strategy is similar to the ones involved in retaining verbatim memory of sentences and memory of semantic content. If it is, we should expect our awareness of direct causation to be malleable. In particular, we would expect that if the experiments were re-run, but with the resultative sentences embedded in some sort of context which made it plausible to look for direct causation, then subjects would fare better at finding it. For instance, there is a well-known connection between direct causation and our intuitions about moral and legal responsibility (Hart & Honoré, 1959; Pietroski, 1998). If resultative sentences were embedded in a context that encouraged subjects to consider whether subjects were morally or legally responsible for the action denoted by the secondary predicate, we should expect them to be much better at identifying the direct causation in the resultative. I have not tested this prediction, but I have presented a few of my non-linguistically trained colleagues with some relevant examples. They all produced appropriate judgements in discourses like: *Mom said Bill would get \$5 if he wiped the kitchen table clean. But she didn't pay him because he just started wiping the table when Jenny was around, and that made her want to do it too, so she finished the job.* (The embedding of resultatives in the manner of Tests 3 and 4 may be another case of this, too.) This prediction is noteworthy because it raises some interesting questions about the relation of the present work and some popular research in evolutionary psychology regarding cognitive modules. In particular, the present work may be relevant to those who posit a “cheater detector” module to explain subjects’ improved performance in certain reasoning tasks when the tasks are embedded in contexts where legal or moral responsibility is relevant (e.g., Cosmides & Tooby, 1992). If we are good at spotting direct causation in contexts where legal or moral responsibility turns on it, and we aren’t so good at spotting it in similar contexts where no such responsibility is implied, then this might constitute some evidence that cheater-detection is playing a role.

In sum, both the general terrain of cognitive science and the more local area of psycholinguistics strongly suggest that our ignorance of some aspects of sentence meaning is unsurprising. In fact, this claim is consonant with other previously established results, and it makes several interesting predictions, some of which have received confirmation. Obviously, a great deal of work needs to be done before any of these speculations are confirmed. (As I mentioned earlier, my goal here is more to raise some questions about language understanding than to answer them.) But in any case, notice that none of these predictions, or the proposal that generated them, must be correct in order for the present point to have been made. The present point was only that it is not at all odd to think that we don’t always understand all of the meaning of a sentence. I contend that we have amply justified that point.

4.2 We are only tacitly aware of syntax

The second conclusion we can draw from the experimental data is that our awareness of some syntactic features is only tacit. Although subjects' awareness of semantics varied from test to test, their awareness of syntax did not. Regardless of the type of priming used, the proportions of correct answers in the syntactic condition were always very low, and they did not differ significantly from one another ($N = 60$, χ^2 (2 df) = .36, $p > .83$). Furthermore, the rate of positive answers to the target sentence in the syntactic condition was extremely close to the rate of positive answers to the filler sentences in which there were no traces. Thus, the few correct positive responses to the target sentences appear to be due to effects other than syntactic knowledge. (This last claim also holds for Test 1, since correction for guessing is known to be a conservative measure.) Since syntactic knowledge is often held to be a paradigm of tacit knowledge, the claim that our awareness of syntax is tacit is to be expected. Indeed, if we determine the syntactic phenomena to be anything remotely like all the major contemporary theories treat it as being, the tacitness of much syntactic structure is almost guaranteed. After all, on this conception of syntactic phenomena, there are a great many intuitively distinct syntactic phenomena. A theory that had a separate rule or property or "explanation" for each of these nominally distinct phenomena would scarcely be a theory at all. Rather, it would be at best a convenient catalog of these phenomena. On the other hand, a theory that derives these phenomena from deeper underlying principles, showing that many of them are in fact instances of one and the same phenomenon, will of necessity re-deploy structural (or lexical) elements present in the nominally distinct phenomena. And that is just the sort of result we need, as a matter of empirical fact, to demonstrate tacit awareness of syntactic structure.

4.3 A sense in which we are aware of sentence-meanings

I argued in Sect. 4.1 that there is an important sense in which we lack an accessible awareness of some semantic features of our sentences. Now I want to turn the tables and argue that there is also an important sense in which we *are* accessibly aware of these features. These two claims are consistent because, as we've seen repeatedly, there's no one tacit/accessible distinction; there are lots of them (and lots of these are useful and important). I'll begin by articulating a rather generous notion of accessibility and show it is useful for thinking about language. Then I'll argue that the priming used in Tests 3 and 4 counts as legitimate for such a tacit/accessible distinction.

Earlier we saw that our awareness of direct causation is tacit in the way that e.g., our awareness of the impact of pupil size on our assessments of beauty is (cf. e.g., the contrasts between Tests 4 and 5). But our awareness of the impact of enlarged pupils is tacit only in a weak sense, because we can become explicitly aware of it with only a little bit of priming. However, the case still counts as tacit awareness (in one sense) because with no priming to help us, we tend not to notice how pupil size affects us. Thus, tests for the tacitness of our awareness may require carefully restricting the type of priming subjects are given. But as we saw in Sect. 1, other cognitive phenomena can be usefully explored by generating explicit awareness with the help of a great deal of priming. In the case of blindsight for example, simply forcing a person to guess what is displayed in the blinded portion of their visual field results in disappointing responses, suggesting that they do not perceive anything there (e.g., Marcel, 1993; Vision, 1998). The interesting result—that persons with blindsight do appear

to receive some perceptual information from their blindsighted region—appears only when subjects are given enough information to reduce their task to identifying which of a few possibilities is correct. Thus, these latter tests require a very liberal type of priming. Just as the first type of test might not show much if a lot of priming is given (because the test would be too easy), the second type of test might not show much if only a little priming is given (because the test would be too hard). These two different attitudes toward priming suggest that the two sorts of tests are measuring different types of cognitive phenomena. Priming is heavily restricted in the first sort of test because it measures the information subjects typically consciously perceive and use in their everyday experience. That is, the first sort of test explores the information we typically use in our daily lives. There is a lot of priming in the second sort of test because it measures information that subjects have *any ability at all* to extract from the environment and use in conscious, rational deliberation. (Intuitively, blindsighted people have some ability to use information about spatial orientation; in contrast, no amount of priming will enable anybody to extract information about whether there are 10 or 20 gold atoms in a given region without using any special equipment.) It is tempting to say that the second sort of test discriminates between those features of the environment that we simply cannot become accessibly aware of without lengthy, explicit formal training, from those features that we can. E.g., the fact that subjects as a group never seemed to recognize syntactic traces despite the large amounts of priming used in the various tests suggests that we do not have any capacity to perceive these elements. However, there is more to be said here, since there is no clear distinction between priming and “lengthy, explicit formal training”. I will touch on this matter below; but for now I simply assume that there is an important cognitive distinction indicated by the amount of priming needed to see direct causation, and the amount needed to get them to see traces.⁹ (Notice also that the distinction I’ve gestured at has an “empirical interpretation” in terms of information we typically use and information we don’t normally use, but in principle can exploit.)

We’ve seen that there are two relevant notions of accessibility—one which requires very little priming to test for it, and one which allows a great deal of priming. (Actually, there are lots of relevant distinctions, but I’ll focus on just these two.) According to the second notion, our awareness of the phonological properties of the letter *t* is probably accessible, since a little prompting is probably enough to get subjects to hear the difference between the *ts* in *stab* and *Batman*. (In contrast, we are probably only tacitly aware of these aspects of phonology according to the first notion.) However, blindsighted persons’ awareness of whether the bar is horizontal or vertical appears to be only tacit. The fact that their ‘guesses’ are significantly better than chance only shows that they had some kind of awareness of the bar’s orientation (e.g., Marcel, 1993; Vision, 1998). They also tend to declare that they didn’t see anything and that they were simply guessing about the bar’s orientation, which suggests that they were unable to use that information in their conscious, rational deliberation. What about the semantic features like direct causation? The evidence we’ve seen suggests that

⁹ Note that the cognitive distinctions just assumed here do not entail a non-vague distinction between priming and lengthy formal training. There can be relevant distinctions among continuous distributions. E.g., for many practical purposes, the range of a logistic function (which produces an S-shaped curve) has three parts: the lower part, the ascending part, and the upper part. Also, I stress that I have not offered a philosophical analysis of a particular tacit/accessible distinction. Rather than analyzing such a notion, I am merely attempting to use it. I leave its proper philosophical analysis for other work.

our awareness of these features is accessible according to the second notion of accessibility. In Tests 3 and 4, subjects were told to indicate that the feature was not present if they felt that they would just be guessing. The high rate of successful responses suggests that they were not only aware of the correct answers, but that they were aware that they were aware, and that they could consciously and rationally manipulate this information. (Recall also that a correct score required correct answers on both the target-yes and target-no questions, so subjects needed a firm enough grasp of direct causation to perceive its absence as well as its presence.) Although there may have been some amount of residual guessing involved, it is unlikely that this effect on the scores would be enough to undermine the main claim. (E.g., even if there was a large residual guessing effect, say 20%, and even if we round up to the nearest whole number to ensure the greatest possible effect, the resulting version of Test 4 would still be significant ($\chi^2(1 \text{ df}) = 5.95, p = .0147$.) (Obviously, subjects were given quite a bit of priming in Tests 1 and 2, as well; thus, the sense in which they have accessible awareness of the target semantic features must be understood to allow for a generous amount of priming.)

Even though the second notion of accessibility is more generous, it is still crucial that subjects were not given too much priming. Were they? Crucially, the priming must not simply give the subject the answers to the target questions. (In the general case, there might be problems about whether a type of priming taught *some* of the answers to the target questions, but in the present case, that does not appear to be so.) Clearly, the strongest form of priming never gave away the answers, primarily because the discussion of direct causation always involved a different sort of linguistic construction than the forms used in the test questions. Moreover, the corresponding form of priming didn't appear to affect the responses in the syntactic condition. To the extent that we are antecedently confident that our awareness of the location of the target trace in *Who do you think Melanie likes?* is only tacit, we can treat the syntactic conditions as measures of whether we have given too much priming. If subjects were given so much priming that they were able to succeed in the syntactic condition, then that would be evidence that an illegitimate amount of priming had been used. The fact that subjects never had much success in this condition provides some evidence that they did not receive too much priming. Of course, this argument only shows that the amount of priming used failed to meet one standard for being excessive. Nonetheless, the differing behavior between the syntactic and semantic conditions is highly suggestive. In each test, the fraction of correct scores in the syntactic condition was consistently smaller than the corresponding fraction for the semantic condition. Moreover, the fractions for the syntactic conditions never changed significantly across any of the tests, while the fractions for the semantic conditions varied quite dramatically. Thus, it is plausible that subjects were consciously sensitive to direct causation in Tests 3 and 4.

Finally, a comment is in order about our focus on our grasp of individual sentences, instead of the traditional focus on our knowledge of a grammar. This shift in attention is relevant from the standpoint of contemporary theoretical linguistics. There is currently a strong trend in linguistics towards assuming that (at least some of) the grammatical system is quite small, consisting of only a few number of computational properties of sentences, while sentences themselves contain a lot of structure (Bresnan, 1982; Chomsky, 1981; 1995, 2000, 2001; Falk, 2001). For example, instead of employing a set of rules that explains why there can be empty categories in some locations but not others—cf. *Jill wants to kiss Bill* (meaning Jill wants Jill to kiss Bill)

versus **Jill wants Bill to kiss* (meaning Jill wants Bill to kiss Jill)—current research is focused on deriving these behaviors from just a few highly abstract properties of syntactic organization (e.g., Hornstein, 2003). In practice, this strategy has tended to increase the amount of postulated structure in sentences. Moreover, much of this sentential structure is lexically specified, as part of the structure of individual lexical items (roughly, words). Thus, intuitively speaking, there is a tendency to let the sentences themselves carry much of the theoretical machinery that drives linguistic theories, leaving only a few very general rules to derive sentences from collections of words. That is, once you have what the lexicon supplies to the sentence, you don't need to know very much more. You'll need to know about a few general syntactic properties, but you won't need to know a long list of detailed, highly specific *rules* that are themselves not part of the sentences. (Indeed, sentential structure has become substantial enough that even the very notion of a sentence has become somewhat intricate.) So attending to sentences may enable us to address more of the structure of linguistic theory than will attention to the grammatical system.

5 Directions for further research

The present project of understanding our cognitive relations to the syntax and semantics of language is at a nascent stage. In the previous sections, I hope to have raised some useful questions about these relations. In this final section I will wrap up with some technical caveats and some philosophically relevant issues. As the section title suggests, I only sketch these caveats/issues, in order to indicate where research in this vein might continue.

I begin with the caveats. First, addressing the issue of our cognitive relations to semantic and syntactic properties of language clearly requires the use of a linguistic theory. Moreover, we have seen that the dependence on linguistic theory is substantial, which means that the details of current linguistic theory, not just its bare outlines, are centrally relevant. Like any other developing scientific field, researchers in linguistics often disagree about the details of particular theories. Thus, the use of a given linguistic theory is a nontrivial feature of the model. Such assumptions appear unavoidable: strong commitments to substantial linguistic details are the norm empirical fields like linguistics, psycholinguistics, and neurolinguistics. (Indeed, although some philosophers object to research that depends on strong commitments to background hypotheses, it is hard to see how the sciences could advance without being able to adopt the “framework” of a body of background assumptions. E.g., some physicists perform experiments based on certain assumptions about dark matter, even though other physicists disagree, and perform experiments based on opposing assumptions.) So the present use of a linguistic theory is only par for the course. As I stressed earlier, to the extent that we take linguistic theories seriously as describing the structure of a particular psychological ability (i.e., our linguistic “competence”), we should be unafraid to explore these assumptions in extralinguistic contexts such as the present one. At the same time, the techniques I have used and the particular properties of syntactic and semantic elements are quite general, and will work across a variety of frameworks. E.g., the results in question transfer straightforwardly to a Minimalist theory that doesn't posit syntactic traces per se, but does have residual markers of the joint operations of Copy and Delete.

Second caveat. The present work is obviously only a very small step towards uncovering our general cognitive relation to semantics. Many more semantically relevant constructions and features need to be studied before we have a clear picture of this phenomenon.¹⁰ The linguistics literature contains many semantic features that appear to be part of natural language: e.g., monotonicity properties, definiteness/indefiniteness, benefactive constructions, stage/individual predicates, etc. (e.g., Larson & Segal, 1995). Yet it is not obvious whether some of these features are more or less accessible than direct causation, and under what priming circumstances. Although I think direct causation provides a good ballpark estimate of our cognitive relation to phonologically null semantic elements, there is obviously much more to be said about the issue. On a related note, we can also raise questions about the circumstances under which subjects' performances will markedly improve or degrade. For instance, since I was generally interested in subjects' abilities to identify linguistic elements within entire sentences, I restricted the amount of focus supplied on the relevant portions of the sentences (cf. below). In particular, subjects were not asked leading questions such as "What happens if you wipe the table clean? Do you (directly) cause the table to become clean?" Informally, such leading questions appear to improve subjects' abilities to identify the relevant construction as expressing direct causation. (Nor should this be surprising; the psychology literature is filled with examples of behavioral studies where minor variants in the presentation of the task significantly affect the outcomes.) This is only one of the many ways that the present experimental paradigm can be manipulated to further explore the question of how and to what degree we grasp the various features of our language. Indeed, from the standpoint of experimental design, many other forms of tests could also be used. Although the present methods generated some interesting results, much more sensitive tests could be used. The psychology literature contains lots of strategies for isolating particular cognitive phenomena and removing extraneous sources of variance. Many of these strategies can be employed to help us better understand what we perceive when we perceive language.

Having made these caveats, I now turn to four philosophically relevant issues.

Issue one: Acquisition versus perception. The first issue concerns the role of priming. Priming raises a question reminiscent of Plato's *Meno*: what is the difference between acquiring a concept and perceiving instances of it (as such)? Initially, this question looks rather odd; surely there's a huge difference between learning what an X is and seeing an X (as an X)! Indeed, for many concepts acquisition occurs without any perception of instances; e.g. we can have the concept of a unicorn, a neutrino, or the number seven without ever seeing any of these things. But in the present case, the issue is different, because a significant amount of priming was needed to get subjects to perceive some semantic feature. As priming becomes increasingly relevant to the ability to spot an X, the difference between learning what Xs are and perceiving Xs becomes increasingly subtle. In the end, there may be no significant difference

¹⁰ Even the feature of direct causation merits further investigation. Although subjects have a certain degree of awareness of direct causation in causative and resultative constructions, how aware are they of this feature in particle constructions such as, *John ate the apple up*? Many linguists think that *up* expresses an end state directly caused by the main action of the verb. In other words, they think that this sentence is a kind of resultative construction, too. But if it is, it is an open question whether it will render direct causation as accessible as it is in ordinary resultatives. It is also open whether there are any appropriate tacit/accessible distinctions that render direct causation accessible in these constructions.

between these two activities, or the border between them may have a large vague penumbra. Other researchers have arrived at similar conclusions; e.g., George Miller and Phillip Johnson-Laird suggest that “any sharp division between conception and perception seems questionable” (1976, 41; for similar sentiments, cf. Chomsky, 1965, 51; Prinz, 2002 explores a related theme). Alternatively, some instances of apparently tacit awareness of Xs may really be circumstances where subjects have no concept of Xs, but are disposed to quickly and easily acquire the concept (cf. Maibom, 2003 for some relevant discussion). I believe, although I cannot argue for it here, that a really clear account of tacit knowledge and tacit awareness (and of concepts and beliefs) will show that many instances of these cases can be *identified* with a readiness to acquire certain concepts or beliefs. But these issues, like all the ones I list below, are matters for further research.

Second issue: Semantic theories and trivial inferences. Many philosophers hold that one understands a sentence only if one can draw all the ‘trivial’ inferences from it (e.g., Harman, 1973, 1993; Peacocke, 1992; Sellars, 1974; Wittgenstein, 1953). Indeed, it is sometimes said that such inferences must be entirely obvious to the speaker. How, for instance, could you be said to understand (10a) without immediately understanding that it entails (10b)?

- (10) a. Mary poisoned the flower and afterwards it wilted.
b. ∴ Something happened after something else.

Moreover, many philosophers hold the stronger view that it is *constitutive* of the meaning of an expression that speakers be able to immediately draw all such inferences from sentences containing it (cf. Fodor & Lepore, 1992 for discussion). So if typical speakers don’t immediately see that *Mary poisoned the flower and Φ it wilted* entails *Something happened after something else*, then by definition, Φ can’t mean what *afterwards* means. E.g., according to Peacocke (1992), the possession of a concept requires that one finds certain inferences “primitively compelling”. Such philosophical views about the nature of meaning make strong empirical claims that are hard to reconcile with what we’ve seen so far. For instance, using a standard logico-semantic apparatus suggests that the inference in (11) is about as trivial as possible.

- (11) a. Mary wiped the table clean.
b. Direct causation exists.

The relevant logic of the inference here appears to be:

- (12) a. $(\exists x)(\text{wipe}(x, M, \text{the table}) \ \& \ (\exists y)((\text{clean}(y, \text{the table}) \ \& \ \text{directly-cause}(x, y))))$
b. ∴ $(\exists x)(\exists y) \ \text{directly-cause}(x, y)$

(12a) is read as ‘There is an event x of Mary’s wiping the table and there is an event y of the table’s being clean and x directly caused y’. Thus, the inference in (11) involves nothing more than simple conjunction reduction, just like the inference in (10):

- (13) a. $(\exists x)(\text{poison}(x, M, \text{the flower}) \ \& \ (\exists y)((\text{wilt}(y, \text{the flower}) \ \& \ \text{after}(x, y))))$
b. ∴ $(\exists x)(\exists y) \ \text{after}(x, y)$

But by looking for instances of direct causation, subjects were attempting to verify the inference in (11). The fact that they often failed despite a good deal of priming

makes it unclear how immediately grasping trivial inferences constitutes an epistemic or metaphysical constraint on sentence meanings. After all, if a person doesn't immediately grasp the inference in (11), it doesn't follow that she fails to understand (11a) in the same way that someone who fails to grasp the inference in (10) doesn't understand (10a). (Although we have seen that there's a sense in which we don't understand some of our sentences, that sense is not the same as simply being incompetent with respect to a sentence.) With respect to the metaphysical claims about what constitutes meaning, it looks like we have a clear case of *modus tollens*. (11a) *does* express direct causation, and the fact that speaker miss this shows that meanings are not constituted by the trivial inferences we draw from them. Instead we should assume that our best linguistic theories tell us what sentences mean. Ordinary speakers' unreflective judgments about the meaning of these sentences constitute only one source of evidence for the linguistic theory. Sometimes these judgments will be overridden. If this is right, then the use of trivial inferences to characterize the epistemology or ontology of language is mistaken, and the associated theories of holism and inferential role semantics are false. The problem for such views of semantics can be seen in the form of a simple argument: (i) The linguistic facts show that expressions of type E mean that M. (ii) If semantic theory T is correct, then E means M only if the inference X is primitively compelling (perhaps because X is relevantly similar to paradigm examples of primitive compellingness). (iii) But, as an empirical matter of fact, X is not primitively compelling (within a relevant population). (iv) Hence, semantic theory T is not correct. Although there's much more to be said about this argument, the problem it exposes is quite general. As soon as you start taking linguistic theories seriously, you start encountering all sorts of features, semantic and otherwise, that typically don't appear on a cursory examination of English. These features are often sufficiently *recherché* that our cognitive relation to them is a topic for empirical investigation. But to the extent that we must investigate our relations to such features, global semantic views like holism or inferential role semantics entail highly nontrivial empirical claims which, I have argued in this paper, are sometimes false.

Of course, there are many ways to try to avoid these conclusions. For instance, one might argue that the inference in (11) is not trivial, and so does not count against the criterion. But to make such a strategy work, we will of course need some independent characterization of what counts as a trivial inference, such that (11) is ruled out, but other inferences like (10) are not. If we just stipulate that (11) doesn't count, then our notion of trivial inference is in danger of being trivialized into a term that—by fiat—applies to just those inferences that speakers find trivial.

Third issue: Semantics versus syntax. The third issue concerns the distinction between syntactic and semantic phenomena. Ordinarily, we think that the difference between syntax and semantics is about as obvious as can be. In a sentence like *John broke the chair*, the word *John* preceding the word *broke* is a syntactic property, and the word *chair's* denoting certain pieces of furniture is a semantic property. How could the two types of linguistic features ever be confused? Such bewilderment comes from assuming an externalist view of semantics. However, when we adopt the internalist view of semantics used in this paper, the syntax/semantics distinction becomes less clear. An internalist view of semantics holds that semantic theories concern how expressions of a language become associated with mental representations of various sorts. That is, an internalist semantic theory only theorizes about how a person's expressions come to be hooked up in the right way to her representational abilities. It is certainly possible that such a theory (or at least its linguistic aspects) can be

articulated just by appealing to the independently motivated morphological and syntactic structure of our expressions. For example, as I mentioned above, it is a common (albeit not universal) view in linguistics to represent direct causation as an element in the total syntactic structure of the sentence. The most accurate way to represent the semantic impact of direct causation may well turn out to be by simply “disquoting” this syntactic structure in a rich metalanguage, such as English with its syntactic structure made explicit. In such a case, (the mental representation of) the syntactic construction expressing direct causation could be identified with (the mental representation of) the meaning of this construction. There will undoubtedly remain many important differences between syntax and semantics, but they may not be relevant to linguistic theory. That is, the distinction between syntax and semantics may not be a linguistic distinction at all, strictly speaking. Rather, it may be a distinction that is better characterized in terms of other cognitive systems that exploit the outputs of the language system. Such a possibility does not vitiate the view sketched in Sect. 4; rather it merely suggests that the issues discussed there may not be linguistic issues after all. Instead, the view may ultimately pertain to some non-linguistic cognitive systems that use language or the outputs of the linguistic faculties.¹¹

Fourth issue: Language understanding and the law. The final issue I will mention concerns the role of language understanding in legal proceedings. Peter Tiersma and Lawrence Solan have shown that the reception of expert linguistic testimony in court cases is mixed: sometimes linguistic testimony is admitted without dispute, but other times it is refused (Tiersma & Solan, 2002). Linguistic experts are often allowed to testify about the meaning of text in a foreign language. However, Tiersma and Solan note that when it comes to expert testimony about ordinary English, things are often much different. Linguistic testimony about the meaning of a discourse has been denied on the grounds that it would not help jurors understand the meaning of the conversation, and might even confuse them. Tiersma and Solan report that the “notion here seems to be: ‘Why do we need a language expert? We have twelve jurors who all speak English. Let them decide’” (233). Similarly, the authors discuss the case of James Free, who challenged his death sentence from a murder trial, on the grounds that the court’s instructions to the jury were misleading and obscure. Here, Tiersma and Solan report ‘The trial court listened carefully to various experts, including a linguist, and decided that jurors likely did not understand instructions on how to decide whether to impose the death penalty. But the court of appeals had an almost disparaging attitude towards the research and quickly affirmed Free’s death sentence’ (237). This dismissive attitude toward linguistic testimony is a mistake. We’ve seen that subjects can have great difficulties achieving explicit awareness of the meanings of relatively short, simple sentences. Their difficulties seem to increase when the relevant component of meaning is not made salient by the discourse. But these “less salient” readings of various sentences could easily be crucial to a court case. How risky is the current legal trend of disallowing expert linguistic testimony? This question could be explored by studying some of the cases where linguistic testimony was relevant. Are there cases where relevant interpretations of text went unnoticed? If so, how many? This general issue can even be explored hypothetically, by examining the plausibility of cases occurring which involve nonsalient interpretations that would go unnoticed without the aid of a linguist. (These hypothetical cases would not be purely speculative, since

¹¹ It would be interesting to compare this issue with the Whorf-Sapir hypothesis that thought is somehow intrinsically dependent upon language (e.g., Bloom & Keil, 2001; Boroditsky, 2001; Fodor, 2001; Higginbotham, 1994).

for any given sentence, it is a straightforward empirical issue of how likely it is that a subject from a given population will notice a given interpretation of the sentence, given some form of background priming.)

6 Conclusion

This paper contains two main claims. First, there is no one epistemic dividing line that separates the tacit from the accessible. Instead, there are many tacit/accessible distinctions, varying according to the sorts of priming that are considered legitimate. Thus, any time a tacit/accessible distinction is used in research, it must be defended as the right distinction to employ for whatever project is at hand. Similarly, a claim that some bit of information is only tacitly known or perceived is always relative to some type of background priming. Second, subjects can be made aware of semantic features like direct causation, but not without a lot of priming. Moreover, their awareness of such features is quite malleable, with significant effects arising variations in the priming. Having motivated two types of tacit/accessible distinctions, we saw that according to one of them, we are merely tacitly aware of some of the smallish semantic features of sentences, but according to the other, we have an accessible awareness of them. In contrast, however, awareness of the target syntactic feature was robustly tacit. Finally, the issue of our awareness of syntactic and semantic features was shown to have implications for a variety of areas. The broad range of these areas—linguistics, psychology, philosophy, and the law—suggests that the present approach to our knowledge of language is a fruitful one.

References

- Anderson, R., & Pichert, J. (1978). Recall of previously unrecalled information following a shift in perspective. *Journal of Verbal Learning and Verbal Behavior*, 17, 1–12.
- Baker, M. (1988). *Incorporation*. Chicago: University of Chicago Press.
- Baker, M. (2003). *Lexical categories*. Cambridge: CUP.
- Bittner, M. (1999). Concealed causatives. *Natural Language Semantics*, 7, 1–78.
- Block, N. (1995). On a confusion about a function of consciousness. *Behavioral and Brain Sciences*, 18, 227–287.
- Bloom, P., & Keil, F. (2001). Thinking through language. *Linguistics and Philosophy*, 16, 351–367.
- Boroditsky, L. (2001). Does language shape thought? Mandarin and English speakers conceptions of time. *Cognitive Psychology*, 43, 1–22.
- Bresnan J.J. (1982). *The mental representation of grammatical relations*. Cambridge, MA: MIT Press.
- Chomsky, N. (1965). *Aspects of the theory of syntax*. Cambridge, MA: MIT Press.
- Chomsky, N. (1981). *Lectures on government and binding*. Dordrecht: Foris.
- Chomsky, N. (1995). *The minimalist program*. Cambridge, MA: MIT Press.
- Chomsky, N. (2000). Minimalist inquiries: The framework. In R. Martin et al. (Eds.), *Step by step: essays on minimalist syntax in honor of Howard Lasnik* (pp. 89–155). Cambridge, MA: MIT Press.
- Chomsky, N. (2001). Derivation by phase. In M. Kenstowicz (Ed.), *Ken Hale* (pp. 1–52). Cambridge, MA: MIT Press.
- Comrie, B. (1985). Causative verb formation and other verb-deriving morphology. In T. Shopen (Ed.), *Language typology and syntactic description*, Vol. III (pp. 309–348). Cambridge: CUP.
- Connolly, T., Arkes, H., & Hammond, K. (2000). *Judgment and decision making*. Cambridge: CUP.
- Cosmides L., & Tooby, J. (1992). *The adapted mind*. Oxford: OUP.
- Crimmins, M. (1992). Tacitness and virtual beliefs. *Mind and Language*, 7, 240–263.
- Culicover, P. (1997). *Principles and parameters*. Oxford: OUP.
- Davidson, D. (1967a). Truth and meaning. Reprinted in Davidson 1984, pp. 17–36.
- Davidson, D. (1967b). The logical form of action sentences. Reprinted in Davidson 1980, pp. 105–122.

- Davidson, D. (1967c). Causal relations. Reprinted in Davidson 1980, pp. 149–162.
- Davidson, D. (1970a). Action and reaction. *Inquiry*, 13, 140–148.
- Davidson, D. (1970b). Semantics for natural languages. Reprinted in Davidson 1984, pp. 55–64.
- Davidson, D. (1974). Belief and the basis of meaning. Reprinted in Davidson 1984, pp. 141–154.
- Davidson, D. (1984). *Inquiries into truth and interpretation*. Oxford: Clarendon.
- Dowty, D. (1979). *Word meaning and montague grammar*. Dordrecht: D. Reidel.
- Dummett, M. (1975). What is a theory of meaning? (I), in Dummett 1993, pp. 1–33.
- Dummett, M. (1978). What do I know when I know a language?, in Dummett 1993, pp. 94–105.
- Dummett, M. (1993). *The seas of language*. Oxford: Clarendon.
- Falk, Y. (2001). *Lexical-functional grammar*. Stanford: CSLI.
- Fodor, J. (1998). *In critical condition*. Cambridge MA.: MIT Press.
- Fodor, J. (2001). Language, thought, and compositionality. *Linguistics and philosophy*, 16, 1–15.
- Fodor, J., & Lepore, E. (1992). *Holism: a shopper's guide*. Oxford: Blackwell's.
- George, A. (1989). How not to become confused about linguistics. In A. George (Ed.), *Reflections on Chomsky* (pp. 90–110). Oxford: Blackwell.
- Gilovich, T., Griffin, D., & Kahneman, (2002). *Heuristics and biases: The psychology of intuitive judgement*. Cambridge: CUP.
- Haegeman, L. (1994). *Introduction to government and binding theory* (2nd ed.). Oxford: Blackwell.
- Hale, K., & Keyser, S. (1987). A View from the Middle, Lexicon Project Working Papers 10, Cambridge, MA: Center for Cognitive Science, MIT.
- Hale, K., & Keyser, S. (1993). On argument structure and the lexical expression of syntactic relations. In K. Hale & S. Keyser (Eds.), *The view from building*, 20, (pp.53–109). Cambridge, MA: MIT Press.
- Harman, G. (1973). *Thought*. Princeton: Princeton University Press.
- Harman, G. (1993). Meaning holism defended. In J. Fodor & E. Lepore (Eds.), *Holism: A consumer's update*, Amsterdam: Rodopi.
- Hart, H., & Honoré, A. (1959). *Causation and the law*. Oxford: OUP.
- Hess, E. (1965). Attitude and pupil size. *Scientific American*, 212, 46–54.
- Hess, E. (1975). The role of pupil size in communication. *Scientific American*, 233, 110–119.
- Higginbotham, J. (1985). On semantics. Reprinted in E. Lepore (Ed.), *New directions in semantics* (pp. 1–54). London: Academic Press. 1987.
- Higginbotham, J. (1989). Elucidations of meaning. *Linguistics and Philosophy*, 12, 465–517.
- Higginbotham, J. (1994). Priorities in the philosophy of thought. *Proceedings of the Aristotelian Society*, Supp. vol. 20, 85–106.
- Higginbotham, J. (1998). Conceptual competence. In E. Villanueva (Ed.), *Concepts* (pp. 149–162). Atascadero, CA: Ridgeview Publishing Co.
- Hornstein, N. (2003). On control. In R. Hendrick (Ed.), *Minimalist Syntax* (pp. 6–81). Oxford: Blackwell.
- Johnson, K. (2004a). Tacit belief, semantics, and grammar. *Linguistics and Philosophy*, 27, 57–91.
- Johnson, K. (2004b). From impossible words to conceptual structure: The role of structure and processes in the Lexicon. *Mind and Language*, 19, 334–358.
- Larson, R., & Segal, G. (1995). *Knowledge of Meaning*. Cambridge, MA: MIT Press.
- Levin, B., & Rappaport Hovav, M. (1995). *Unaccusativity: At the syntax-lexical semantics interface*. Cambridge, MA: MIT Press.
- Lewis, D. (1973). Causation. *Journal of Philosophy*, 70, 556–567.
- Loftus, E., & Hoffman, H. (1989). Misinformation and memory, the creation of new memories. *Journal of Experimental Psychology: General*, 118, 100–104.
- Loftus, E., & Ketcham, K. (1994). *The myth of repressed memory*. NY: St. Martin's Press.
- Maibom, H. (2003). The mindreader and the scientist. *Mind and Language*, 18, 296–315.
- Maier, N. (1931). Reasoning in Humans, II. *Journal of Comparative Psychology*, 12, 181–194.
- Marcel, A. (1989). Phenomenal experience and functionalism. In A. J. Marcel & Bisiach E. (Eds.), *Consciousness in contemporary science* (pp. 121–158). Oxford: OUP.
- Marcel, A. (1993). Slippage in the unity of consciousness. In G. Bock & J. Marsh (Eds.), *Experimental and Theoretical Studies of Consciousness* (pp. 283–300). Chichester: Wiley.
- Miller, G., & Johnson-Laird, P. (1976). *Language and perception*. Cambridge: CUP.
- Murphy, G., & Shapiro, A. (1994). Forgetting of Verbatim Information in Discourse. *Memory and Cognition*, 22, 85–94.
- Nagel, T. (1969) Linguistics and epistemology. In S. Hook (Ed.), *Language and philosophy* (pp. 171–182). New York: NYU Press.
- Nisbett, R., & DeCamp Wilson, T. (1977). Telling more than we can know: verbal reports on mental processes. *Psychological Review*, 84, 231–259.

- Owens, J., Bower, G., & Black, J. (1979). The 'Soap Opera' effect in story recall. *Memory and Cognition*, 7, 185–191.
- Parsons, T. (1990). *Events in the semantics of english*. Cambridge, MA: MIT Press.
- Peacocke, C. (1975). Truth definitions and actual languages. In G. Evans & J. McDowell (Eds.), *Truth and meaning* (pp. 162–188). Oxford: OUP.
- Peacocke, C. (1992). *A study of concepts*. Cambridge, MA: IT Press.
- Pietroski, P. (1998). Actions, adjuncts and agency. *Mind*, 107, 73–111.
- Pietroski, P. (2000). *Causing actions*. Oxford: OUP.
- Plaut, D. (2003). Connectionist modeling of language: Examples and implications. In M. Banich & M. Mack (Eds.), *Mind, brain, and language* (pp. 143–167). Mahwah, NJ: Lawrence Erlbaum Associates.
- Prinz, J. (2002). *Furnishing the mind*. Cambridge, MA: MIT Press.
- Pustejovsky, J. (1991). The syntax of event structure. *Cognition*, 41, 47–81.
- Sachs, J. (1967). Recognition memory for syntactic and semantic aspects of connected discourse. *Perception and Psychophysics*, 2, 437–442.
- Schacter, D. (1989). On the relation between memory and consciousness: dissociable interactions and conscious experience. In H. L. Roediger III & F. I. M. Craik (Eds.), *Varieties of memory and consciousness: essays in honor of Endel Tulving* (pp. 355–389). Hillsdale, NJ: Erlbaum.
- Sellars, W. (1974). Meaning as functional classification. *Synthese*, 27, 417–437.
- Simons, D., & Chabris, C. (1999). Gorillas in our midst: sustained inattention blindness for dynamic events. *Perception*, 28, 1059–1074.
- Stich, S. (1978). Beliefs and subdoxastic states. *Philosophy of Science*, 45, 499–518.
- Tenny, C. (1994). *Aspectual roles and the syntax-semantics interface*. Dordrecht: Kluwer.
- Tiersma, P., & Solan, L. (2002). The Linguist on the witness stand: forensic linguistics in american courts. *Language*, 78, 221–239.
- Trivers, R. (1985). *Social evolution*. Reading, MA.: Benjamin/Cummings.
- Vision, G. (1998) Blindsight and philosophy. *Philosophical Psychology*, 11, 137–159.
- Wittgenstein, L. (1953). *Philosophical Investigations* (G. E. M. Anscombe, transl.). New York: Macmillan and Co.
- Wright, C. (1986). Theories of meaning and speaker's knowledge. In *Realism, meaning and truth* (2nd ed., pp. 204–238). Oxford: Blackwell, 1993.