Fly swatting: Davidsonian truth theories and context

R. M. Sainsbury

University of Texas at Austin and King's College London

Abstract: A Davidsonian truth-theorist might worry about context-dependence for two kinds of reason. On the one hand, it raises fundamental doubts about the Davidsonian project, which seems to involve isolating specifically semantic knowledge from any other knowledge or skill in a way reflected by the ideal of homophony. Indexicality forces a departure from this ideal, and so from the aspiration of deriving the truth conditions of an arbitrary utterance on the basis simply of axioms which could hope to represent purely semantic knowledge. In defence of Davidson, I argue that once his original idea for dealing with the familiar indexical expressions is suitably implemented, in terms of theorems expressing conditional truth-conditions, the contrast between semantic and non-semantic knowledge occurs in an appropriate place: between knowledge of the conditional truth conditions themselves (semantic), and knowledge of their antecedents (non-semantic). For example, a simplified version of a conditional truth condition for an utterance *u* of "That is a cat" will say that for all *x*, if the utterer in uttering "that" in *u* thereby referred to *x*, *u* is true iff *x* is a cat. The conditional fact belongs to semantic knowledge; non-semantic knowledge is required to work out what the utterer referred to; the two kinds of knowledge conspire to enable a truth condition to be detached.

The other kind of worry relates to difficult technical questions, relating to specific idioms, raised by context-dependence. I consider whether an argument by Travis (2006), designed to establish that there can be no correct truth conditional semantics, can be turned on its head, and used to establish a general method for bringing all forms of context-dependence within Davidson's framework (using quantification into the qualifier "on understanding *U*"). I argue that this will not work, and that we are committed to a piecemeal examination of various cases, of which I give some examples. One of these relates directly to Travis's argument concerning the context-dependence of "grunt". In addition, Davidsonians need to have a sharp eye for the distinction between pragmatic and semantic content, to be specified by a test akin to Grice's notion of cancelability; they should treat sentences like "It's raining" by the conditional truth-condition method; and they should treat possessives as semantically very unspecific. It is part of the thesis of the paper that this list of examples by no means exhausts the cases a Davidsonian needs to address.

1. The ideal of homophony

One key idea in Davidson's approach to language is that the output of a semantic theory should focus on *interpretation*. To interpret an utterance is to redescribe it in a way that will give an ideally correct and complete answer to the question "What did so-and-so say?", asked by an ordinary language user in a typical context. Normally this redescription will be one that the utterer herself should and can accept, and so will not involve concepts (for example the concept of a set or infinitary sequence) which need not be exercised in competent use. This immediately distances Davidson's project from one which takes as central the provision of revealing descriptions of what individual words mean. No doubt for some purposes it is useful to say of the word "snow" that it refers to atmospheric water vapor frozen into ice crystals and falling in light white flakes or lying on the ground as a white layer. But one who says "Snow is white" does not say that atmospheric water vapor frozen into ice crystals and falling in light white flakes or lying on the ground as a white layer is white; he says that snow is white. Davidson's project val-

orizes "homophonic" interpretations, ones in which the very same words are used in the interpretation as were used in the utterance to be interpreted. Is this is a missed opportunity? We should recall, first, that all languages can be expected to have primitive expressions, ones whose meaning cannot be specified by using others (color words are obvious examples); and, second, that Davidsonian interpretations have to be derived from word-related axioms "by logic alone", a demand which provides a quite specific perspective on what it is to reveal the compositional mechanisms of a language.

Homophony is sufficient but not necessary for interpretiveness. Its sufficiency is illustrated by the fact that an utterer of "Snow is white" cannot demur at being reported as having said that snow is white. Its lack of necessity is illustrated by the fact that one can correctly report an utterance using a language the utterer does not understand, thereby describing her utterance in a way she would not recognize as correct. A monolingual Italian who says "La neve e bianca" thereby says that snow is white, even if she cannot recognize as correct this non-homophonic description of what she said.

Indexicality itself shows that homophony is not necessary for interpretiveness. If you utter "I am hungry" you do not say that I am hungry but that you are. Interpretation seems to bring together two kinds of knowledge. There is standing knowledge about the "meaning" of "I", and occasion-specific knowledge of who uttered the sentence. Both are required for interpretation, though the latter kind of knowledge is not the property of a special language module. It belongs to quite general cognitive abilities (as is perhaps more conspicuous in the case of "that", where all kinds of information can be used in discovering the referent). This means that, assuming Davidsonian theory delivers a full semantic account, semantic information is not sufficient for interpretation. The upshot is that Davidson's project is doomed, if it is construed as using truth theory to organize information knowledge of which would suffice for interpretation. But we can broaden the Davidsonian perspective: think of semantic knowledge as sufficient for interpretation only in combination with non-semantic knowledge. We still might be able to use Davidson's approach both to constrain what counts as interpretation, and to do full justice to the specifically semantic aspects of language use. That is the optimistic idea I will pursue in this paper.

2. Classical indexicality

Davidson said that indexicality is a "very large fly in the ointment" (1967: 33), but he made some preliminary suggestions about how it might be handled truth-theoretically. Here are two:

> 'I am tired' is true as (potentially) spoken by p at t if and only if p is tired at t.

> 'That book was stolen' is true as (potentially) spoken by p at t if and only if the book demonstrated by p at t is stolen prior to t. (1967: 34)

These are presumably schemata designed to indicate the form that appropriate T-theorems should take. They do not as such indicate the axioms from which such theorems might be derived. We cannot accept the schemata. An instance of the second is:

> "That book was stolen" is true as (potentially) spoken by Davidson at noon on 02/02/02 if and only if the book demonstrated by Davidson at noon on 02/02/02 is stolen prior to noon on 02/02/02. (Sainsbury 2005: 54)

This is not interpretive. The report draws on concepts, like the time and date, which an utterer of the target sentence would have no need to exercise in a fully competent utterance.

We start to address these problems, while also ensuring a clear demarcation of the province of semantics, by thinking of our goal as the provision of conditional rather than absolute truth conditions, as envisioned by Higginbotham, for example
> If u is an utterance of "today is July 4" by s, and s refers with the utterance of "today" therein to δ, then
> u is true ↔ δ is July 4. (Higginbotham 1994, p. 94)

We still have uninterpretive results. Start with this instance of Higginbotham's schema:
> If u is an utterance of "today is July 4" by Sally, and she refers with the utterance of "today" therein to July 3, then
> u is true ↔ July 3 is July 4.

If we are to use this for interpretation, we will presumably use the antecedent to derive:
> In uttering "today is July 4", Sally spoke truly iff July 3 is July 4

and then report Sally as having said that July 3 is July 4. But one who says this is manifestly irrational, whereas one who uttered what Sally uttered may not be. We therefore need to refine Higginbotham's suggestion, in terms of what I have called the scene/content structure. Intuitively, we want to report Sally along these lines:
> On July 3, Sally said that it was July 4.

The "it" depends anaphorically upon the interpreter's use of "July 3" outside the content reported. The report is entirely neutral about how Sally referred to July 3, and so is consistent with her having uttered something expressing a belief she could rationally hold. We could implement this by replacing the second occurrence of δ in Higginbotham's schema by "it" (perhaps with subscript δ). Then, applying obvious background information, we might end up with a redescription of Sally's utterance as follows:
> Sally uttered "today is July 4", and by her use of "today" in that utterance referred to July 3. So what she said is true iff it [July 3] is July 4.

This is near enough to the intuitively desired form of report. This approach relinquishes the idea that I can always make myself an absolute samesayer with someone whose speech I correctly report. I may have to do some scene-setting (e.g. "Sally, speaking on July 3, said that …") before I can introduce the content-specifying part ("it was July 4"), and this is not a stand-alone specification of a content, but rather one that essentially looks backwards, through anaphoric dependence, to the scene-setting.

A Davidsonian supplies user-instructions along with his truth theory. In the unreformed version, they go something like this: if you encounter an utterance u of a sentence s by a speaker a, check the truth theory for a canonical theorem for s; if it says that s is true iff p, conclude that a said that p. Reforming these instructions for the proposed revision of Davidsonian truth theory, they become rather more complex, on the following lines:

1. Derive a conditional truth condition for s. It will be along the lines: "If in an utterance u of s the speaker referred to objects x, y, and z then u is true iff …x …y … z….

2. Determine each of the referents, and select corresponding expressions x*, y* and z* (constants, not variables) to refer to them.

3. Take an instance of the conditional truth condition in which the starred constants replace the corresponding variables in the antecedent, and suitably related anaphoric pronouns replace them in the right hand side of the biconditional consequent. The result will be along these lines: "If in an utterance u of s the speaker referred to objects x*, y*, and z* then u is true iff …it_{x*} …it_{y*} … it_{z*}….

4.Delete the "If", the left hand side of the biconditional consequent, and the biconditional sign ("iff"); replace the "then" by "he/she thereby said that". The result will have the form: In her utterance u of s referring to x*, y* and z* she thereby said that …it_{x*} …it_{y*} … it_{z*}….
If all goes well, this should be an interpretive report.

The conclusion is that the spirit of Davidsonian truth theory can accommodate at least these kinds of cases of indexicality, and in the process makes an intuitively plausible division between genuinely semantic knowledge (represented by the conditional truth conditions) and the collateral knowledge needed to attain interpretations.

3. Radical Contextualism against truth conditional semantics.

Charles Travis has made the bold claim that no sentence completely expresses a thought, so no sentence's content can be captured truth theoretically: call this Radical Contextualism. He says that it:

> blocks truth conditional semantics. For suppose I say, 'The sentence "Sid grunts" is true iff Sid grunts'. Either I use that last "grunts" on some particular understanding of being a grunter – one understanding among many – or I do not. If I do, then I assign the sentence a property it does not have. For it does not speak of being a grunter on any special understanding of this. But if I do not, then I fail to state any condition under which anything might be true. Being a grunter on no particular understanding of being one is just not a way for Sid to be. (Travis 2006: 47–8)

Sentences containing "grunt" do not completely express a thought, for the words do not settle how "grunt" is to be understood—what sort of behavior is to count as making one a grunter (is one little grunt enough or does one have to be a habitual or serial grunter?). This is resolved in different ways in different contexts. (Compare Travis's examples in earlier work: a brown leaf painted green may count as green in some but not

other contexts.) The first horn of the dilemma for truth theory is that we treat the right side of the biconditional which specifies truth conditions as completed by context, so that it completely expresses the thought that S; then the biconditional will falsely say that every utterance of "Sid grunts" is true iff S. Alternatively, if the right side is not completed by context, we do not have a genuine biconditional: "I fail to state any condition under which anything might be true".

In his early paper (1967), Davidson addressed indexicality by quantifying over persons, times and places. This enabled differences in persons, times and places to lead to different T-theorems. Might one not extend this idea, and deal with context-sensitivity at a single blow by quantifying over contexts? As a first attempt, we might aim at T-sentences on these lines:

> s is true as uttered in C iff: in C, p.

The sentence in the slot marked by "p" translates the sentence referred to by what is in the slot marked by "s". We might initially be encouraged by examples like this:

> "All beer is good" is true as uttered in Australia iff: in Australia, all beer is good.

This may sound true, at least given various simplifying assumptions. But any semblance of a serious contribution to our problem is illusory. The biconditional does not even begin to do proper justice to sensitivity to context. It is derived from a schema which also delivers:

> "All beer is good" is true as uttered in Sydney iff: in Sydney, all beer is good.

Since an utterance in Sydney is also an utterance in Australia, unrestricted application of the schema will yield distinct and potentially conflicting truth conditions for the same unambiguous utterance.

There is a distinct problem, one which confronts an aspect of Borg's recent defence of truth theoretic semantics. Discussing seemingly monadic versions of expressions which also have an explicitly relational form ("ready" [ready for], "married" [married to], "raining" [raining at], "continue" [continue doing such-and-such]) Borg suggests truth conditions on the following lines:

> If u is an utterance of 'Jane can't continue' in a context c then u is true iff Jane can't continue something in c. (2004: 230)

Bill and Sally are talking about Jane; take this fact as an instantiation on the implicitly universally quantified variable c. Bill is trying to communicate to Sally that Jane cannot continue with the marketing research project she started a month earlier. Borg's generalization yields:

If *u* is an utterance of 'Jane can't continue' in a context in which Bill and Sally are talking about Jane then *u* is true iff Jane can't continue something in a context in which Bill and Sally are talking about Jane.

On all likely scenarios, Jane isn't doing anything in the Bill-Sally context (they are talking about her in London, and she is asleep in California), and if she isn't doing anything in it there's nothing she can continue doing in it. So the right side will be false on all likely scenarios, whereas the left side might be true. The problem here is that what we need on the left side is the conversational context, whereas what we need on the right side is some activity which that context determines. It would be surprising were we to get both these different things using two occurrences of a single variable; we certainly could not count on any such coincidence.

If we reconsider Travis's dilemma for truth conditional semantics, it may seem we have material which would lead to just such a surprise. Concerning a putative T-theorem

"Sid grunts" is true iff Sid grunts,

the dilemma was that either the right side is taken in a way divorced from any specific understanding of what it is to grunt, in which case it does not provide a genuine condition at all, or else a specific understanding is in play, in which case it says falsely that this is the only way for "Sid grunts" to be understood. "Understandings", as Travis uses the term, relate both to sentences (like "Sid grunts") and to ways things can be (like being a grunter), so perhaps it is a notion which can effect just the kind of transition needed by Borg's project. (This would be ironic, for Travis was out to attack truth conditional semantics, whereas Borg defends them.) An understanding of a sentence is some kind of additional constraint on what it takes for the sentence to be true. An understanding of a kind of event or state of affairs is an additional constraint on what is involved in its obtaining. (We do not need to insist that this notion of understanding is entirely unequivocal. It is enough that for every understanding of a sentence there is a unique understanding* of a corresponding state of affairs.)

By quantifying over understandings, we might coordinate further specificity for the sentence "Sid grunts" with further specificity concerning what it is to grunt, along these lines:

s, on the understanding *U*, is true iff: on the understanding *U*, *p*.

Let "*s*" be replaced by a name for "Sid grunts", and "*p*" by that sentence itself. Truth is a property of a sentence–understanding pair. There is no absolute what it is to grunt, but only understandings of what it is to grunt. All truth theory need do is coordinate the understandings. Suppose on one understanding of the sentence "Sid grunts" it says that Sid makes grunting noises in the course of the majority of his conversational exchanges in the year 2006, and this (or the correlated) understanding of what it is for Sid to grunt is that it is for him to make grunting noises in the course of the majority of his conversational exchanges in the year 2006. Then an instance of the schema just displayed is:

> "Sid grunts", on the understanding to make grunting noises in the course of the majority of his conversational exchanges in 2006, is true iff: on the understanding to make grunting noises in the course of the majority of his conversational exchanges in 2006, Sid grunts.

If this is intelligible, we seem to have steered between Travis's dilemma. The instance speaks of grunting on an understanding, and so does state a condition under which something can be true. The generalization does not mention any understanding in particular, and so is not open to the charge that it associates a sentence with an excessively specific truth condition, one favoring just one of many specific understandings.

This is all mere hocus pocus. (No doubt Travis would enthusiastically agree.) On a Davidsonian picture, an interpreter is supposed to use the truth theory to arrive at interpretations. But once understandings are quantified over, an interpreter will need to know which understandings are appropriate in order to arrive at an interpretation. This knowledge in itself, however, is the semantic knowledge the theory was supposed to represent, but instead of being represented, it is presupposed.

To fill out this point, we can turn to the presumed axiomatic basis of truth theory. Truth conditional semantics are supposed to be compositional. The quantification over understandings cannot be suddenly imposed upon standard unquantified T-theorems: if these are true, the quantification is unnecessary and inappropriate, and if they are false, the theory is false. We would need to think of axioms applying to expression-understanding pairs, not in an understanding-by-understanding way (for there are too many possible understandings to itemize), but in some general way, for example:

> for all "grunt"-appropriate understandings U, for all x, x satisfies <"grunts", U> iff: on U, x grunts.

The restriction on understandings to ones appropriate to "grunts" is essential. Otherwise we will have instances like:

> for all x, x satisfies <"grunts", the understanding that it is enough to be painted green> iff: on the understanding that it is enough to be painted green, x grunts.

The anomalous understanding is one that would be needed in an account of "green", so it will be a member of the domain of quantification over understandings. But this instance is either nonsense, or it delivers the wrong result. (For the second alternative: an understanding of what it is to be a grunter on which it is enough to be painted green would ensure that suitably painted benches are grunters.) Hence understandings must be restricted to ones appropriate to "grunts". Using the theory now presupposes that the user knows in advance which these understandings are, for the theory does not say. But to know which understandings of "grunts" are appropriate entails knowing what "grunts" means. The knowledge the theory was supposed to state has not been stated but has been presupposed. I conclude that a Davidsonian cannot respond to the kind of global argument we have considered by quantifying over contexts or understandings.

4. Unspecific meanings

How, then, should Travis's point be met? For examples like "grunt" I think the main thing is to distinguish the common phenomenon of unspecific meaning from semantic context-dependence. I'll illustrate with an example that will be uncontroversial for many Contextualists (though I fear may not be so for Radical Contextualists of Travis's kind). There are many ways to run, east or west, to work or to the gym, in the morning or in the evening. An utterance merely of "John runs" does not provide any of these details, though if the utterance is true, it will be made true by an event which resolves every such issue. Context may make some more specific way of running salient, but in doing so, the semantics are not touched. The test is that one can coherently deny that John runs in a salient way, without this being either a retraction or a contradiction. Hence the salient way of running is not part of the semantics.

For example, it would be natural to interpret the utterance of "John runs", as it occurs in the following context

> Jill walks to work. John runs.

as committing the utterer to the claim that John runs to work. The question is whether this commitment (supposing it to be genuine) emerges from the semantics of "John runs". A negative answer is suggested by the following possible variant:

> Jill walks to work. John runs. Indeed, he runs 20 miles a week. But never to work, on account of the traffic.

The coherence, and absence of retraction, suggests that in this utterance the semantics of "John runs" does not assign it the content "John runs to work". This suggests that the same is true of the shorter utterance, for the longer one has the shorter one as a proper part. By the time the interpreter had reached the second full stop in the longer utterance, he should presumably have reached just the state he would have reached when interpreting the shorter one, and so, on the rival view, would have believed that John had been said to run to work. Such an interpreter would have to regard the remainder of the longer utterance as either containing a contradiction or a retraction of the earlier part. Intuitively, however, that is not the case.

The same point can be reached by a slightly different route. Consider

> Bob walks to work. Jill doesn't run. But she runs a quarter marathon every Sunday.

On a Contextualist view, it should be easy to hear this as consistent, for the second sentence will be equivalent to "Jill doesn't run to work". In fact it is hard to hear the whole as consistent, suggesting that the second sentence tells us that Jill doesn't run anywhere (or in any way).

Unspecific meaning is the category to which the Davidsonian should assign Travis's grunter. True, there are many ways of grunting, as there are many ways of doing anything. If an attribution of grunting is true, it is made so by some specific form of grunting. None of this entails that the semantic content of "grunts" varies from context to context (nor that the pragmatic content varies). A test is this: if we can add something equivalent

to "in some way or other" without making a significant difference, the verb is semantically neutral concerning the way it is to be satisfied. The default reading of "Sid grunts" is that he grunts in some way or other: the truth conditions are unspecific relative to various modes of grunting. The default reading of "Sid doesn't grunt" is that he doesn't grunt in any way.

The same goes for color terms: "red" applies to the things that are red in any one of possibly indefinitely many ways (on the inside, on the outside, naturally, through being painted, etc.). If one of these ways is highly salient, we may criticize a speaker for applying "red" to something not red in the salient way; we may voice this criticism by saying that what the speaker said is not true. But we normally do not care about the distinction between semantic content and what a speaker meant, and so we would not discriminate between these different targets of our criticism. Suppose external redness is salient, and someone says, of something which is red inside but not outside, that it is red. It would be natural to respond like this:

> You're wrong: it's not red in the relevant way. It may be red inside, but it's not red outside.

The whole exchange does not require any more specific semantics for "red" than that it is satisfied by something which is red in some way or another. This permits a sensible story about "red inside", which is hard to tell if the salience of external redness made the contained occurrence of "red" apply only to things externally red.

These cases contrast sharply, I believe, with other Contextualist examples. "Jill is ready" is not equivalent to the near-trivial "Jill is ready for something or other", and "This girder is strong enough" is not equivalent to the trivial "This girder is strong enough for something or other". We should not let the fact that there are specific ways to grunt undermine our confidence in the full correctness (barring considerations related to tense) of the claim that "Sid grunts" is true iff Sid grunts. By contrast, "ready" and "enough" may well demand a treatment which reveals their content as context-sensitive.

We cannot offer a global response to the global argument offered by Radical Contextualists. We have to look at their examples case by case. One avenue of response is to say, as with "grunts", that they confuse many ways in which a sentence can be made true with many distinct contents, when really we have a single rather unspecific content. This will not do for all cases. In what follows, I consider further weapons which a Davidsonian should stock in his armory.

5. Other context-driven problems for truth theories: rain and possession

There are very many locutions and kinds of locution that have been held to cause trouble for truth theoretic semantics: small flies in Davidson's ointment (or parts of the one large fly). A thesis of this paper is that there is no general argument to show that they can all be dealt with truth-theoretically. Piecemeal examination is required. In this final section, I will consider just two examples: "It is raining", and possessives.

Suppose that "It's raining" is uttered in Austin on 07/07/07. Let's suppose, with the contextualist, that the semantic content of the utterance is that it's raining in Austin on 07/07/07. (This could be partially justified by the apparent inconsistency of adding "But not anywhere".) A Davidsonian apparently cannot do justice to this as there is no lexical element in the utterance corresponding to a location.

This case can be addressed by exploiting conditional truth conditions. We can treat the sentence uttered as essentially unstructured (tense once more is set aside), so that it receives an axiom along the following lines:

> if in an utterance u of "it's raining" the speaker referred to place or range of places p, then u is true iff it's raining somewhere there [within p].[1]

("There" indicates the use of scene/content structure.) Non-semantic knowledge will be required to determine the range of places. This gives the right result even for cases in which the speaker does not refer to her current location. If my wife in Austin utters

> Mark's in London. It's raining

she would normally be rightly understood as having said that it's raining where Mark is (in London). It also gives the right result for more recherché cases, like the example offered by François Récanati (2002: 317), in which many rain detectors are each monitoring one of many locations, each set to sound a bell if rain is detected. Hearing the bell, but not knowing which detector has sounded, someone exclaims "It's raining". He has referred to the range of places being monitored and has said that it's raining somewhere there.[2]

Turning to possessives, the data are that a noun phrase formed with a possessive (e.g. "John's car") may invoke a wide range of distinct relations. One Davidsonian response would be the conditional truth condition approach:

> for all possession relations R, all objects z, all referring expressions X, all predicates Y:
> if in an utterance u of "X's Y" the speaker refers to relation R, then u thereby refers to z iff z = the satisfier of Y which stands in relation R to the referent of X.

A problem with this approach is that it may deliver something uninterpretive. For example we could use some collateral knowledge, along with the envisaged semantics, to infer that one who uttered

> John's car is rusty

could properly be reported as follows:

> Referring to the car with VIN number 1079856291 he said that it was rusty.

---

[1] Davidson himself, without comment, inserts location into the truth condition:
"'Es regnet' is true-in-German when spoken by x at time t if and only if it is raining near x at t" (1973: 135). The example in the text above shows that this way of incorporating location will not always give the right truth condition.

[2] Rain has been a surprisingly frequently discussed topic recently. For an overview of some positions, see Recanati 2007.

Even if true, this is not interpretive. In contrast to reports of bare demonstratives, some conceptual material relating to cars and John should occur in a fully adequate report.

An alternative approach sees this as a case of deep unspecificity. Some evidence for this might come from the fact that whereas one who utters a sentence containing "John's car" should have some kind of idea what car is at issue, he may have little thought for the relation which holds between it and John. This suggests something like:

"X's Y" refers to the referent of "X"'s satisfier of "Y".

In short, the proposal is that this case be treated like "grunts" and "runs": the possessive construction is simply unspecific in its meaning. One will typically hope that context will reveal greater specificity, but this is not semantically required.

These suggestions indicate some of the weapons a Davidsonian will need to sustain the viability of truth conditional semantics. No doubt there will be others (a firm distinction between the pragmatic and the semantic will be one). And perhaps there are some kinds of contextual phenomena which defeat a Davidsonian approach. As far as I am aware, however, these have yet to be described.

References

Borg, E. (2004). *Minimal Semantics*. Oxford, Clarendon Press.

Davidson, D. (1967). "Truth and meaning." *Synthese* 7, reprinted in his *Essays on Truth and Interpretation*, Oxford University Press, 1984: 17–36.

Davidson, D. (1973). "Radical interpretation." *Dialectica* 27, reprinted in his *Essays on Truth and Interpretation*, Oxford University Press, 1984: 125–39.

Higginbotham, J. (1994). "Priorities in the philosophy of thought." *Supplementary Proceedings of the Aristotelian Society* 68: 85–106.

Recanati, F. (2002). "Unarticulated constituents." *Linguistics and Philosophy* 25: 299–345.

Recanati, F. (2007). "It is raining (somewhere)." *Linguistics and Philosophy* 30: 123–46.

Sainsbury, R. M. (2005). *Reference without Referents*. Oxford, Clarendon Press.

Travis, C. (2006). "Insensitive semantics." *Mind and Language* 21(1): 39–49.