

## *Formal semantics and intentional states*




EMMA BORG


My aim in this note is to address the question of how a context of utterance can figure within a formal, specifically truth-conditional, semantic theory. In particular, I want to explore whether a formal semantic theory could, or should, take the intentional states of a speaker to be relevant in determining the literal meaning of an uttered sentence. The answer I'm going to suggest, contrary to the position of many contemporary formal theorists, is negative. The structure of this note is then as follows: first, I'll very briefly sketch three distinct forms of semantic theory. One, 'strong formal semantics', will be seen to be immediately problematic, leaving us with two other options: use-based theories and what I'll term 'moderate formal semantics'. If we opt for the latter position, the question arises of what kinds of appeals to a context of utterance are legitimate given a formal outlook. I'll suggest that this question arises in two distinct ways and explore the moderate formal semanticist's position in regard to both. However, the conclusion I will reach is that what is characteristic of formal semantics is that it makes only the most minimal semantic concessions to context.

To begin with, if we are thinking about formal approaches to linguistic meaning, it seems that we might adopt what I'll call 'strong formal semantics'. On this kind of account the right point at which to offer semantic analyses is that of formally described sentence-types. Thus literal meaning is held to be entirely context-invariant – a sentence, individuated in terms of its syntax, possesses the very same meaning no matter when, where, or by whom it is produced. Of course, what is communicated by the utterance of a given sentence may change, but the thought is that, as far as *semantics* is concerned, we should abstract entirely from speakers and contexts and concentrate just on the meaning of sentence-types as formal objects of study.<sup>1</sup>

Now, even without spelling out this strong conception of formal semantics in detail, it looks as if it will face problems concerning those sentences whose literal (truth-conditional) meaning apparently depends in some systematic way on the context in which it is produced. Most obviously, we might wonder how to accommodate sentence-types which contain overtly context-sensitive elements (e.g. indexicals, demonstratives, or tense

<sup>1</sup> This sort of position might be associated with the 'null context' approach of Katz 1977.







markers).<sup>2</sup> Somewhat less obviously, we might also be concerned about sentences which seem covertly context-sensitive, containing what are sometimes termed ‘hidden indexicals’ or ‘unarticulated constituents’ (e.g. the comparison class often assumed to be needed to grasp the truth-conditions of a sentence like ‘Sachin is short’ or the location needed to provide a truth-condition for an utterance of ‘It’s raining’) (cf. Recanati 2002). In all these cases, it looks as if a proper specification of meaning can be given *only* once some appeal to the context of utterance has been made.

If we are swayed by this idea that at least some natural language sentences require an appeal to features of the context of utterance in order to provide a specification of their meaning, then strong formal semantics must be rejected. In its place, it seems that two alternative approaches emerge. On the one hand, we might reject the formal approach altogether and opt for a much more context-sensitive, use-based theory of meaning, according to which the right candidates for meaning are utterances, not some more abstract notion of sentence-types. Alternatively, we might seek to preserve the formal approach to the study of meaning but look for some more moderate version of it. On such a more moderate approach it would still be (syntactically individuated) sentence-types which provided the input to a semantic theory, but the context-sensitivity of linguistic meaning would be accommodated by talking of a sentence-type *relativised to a context of utterance*. This kind of approach is well-known from a range of accounts, including Kaplan’s theory of demonstratives (Kaplan 1977). Clearly, however, if we adopt this kind of more moderate formal semantics then there are questions to answer concerning the kind of context-sensitivity which such an approach can allow. In what way, we might ask, can a formal theory of meaning make appeal to a context of utterance and yet still remain a formal account? Kaplan’s technical answer was to treat the context of utterance as an ordered set of parameters against which the sentence-type could be assessed for meaning. Yet, as Kaplan himself noted, this leaves open the issue of the scope and content of these parameters. It is this question, then, that I want to focus on – what kind of appeals to a context of utterance are licensed within a moderate formal semantics?

This issue surfaces in two distinct ways: either we might ask ‘*when* are appeals to the context of utterance admissible in specifications of literal meaning?’ or we might ask ‘*what kind* of features of context of utterance can be appealed to in specifications of literal meaning?’ Considering the

<sup>2</sup> The suggestion is not that strong formal semanticists have nothing to say about the treatment of such items, but rather that whatever they say must fall short of proffering an account of truth-conditional content. Thus, if we expect our semantic theory to deliver truth-conditions, such accounts look problematic.









former, 'when', question, it seems that there are two possible answers. Either appeals to the context of utterance must be limited to those occasions when a syntactic constituent of the sentence demands an appeal to the context of utterance (e.g. the sentence contains a demonstrative or an indexical), or such appeals need not be limited in this way (i.e. not all appeals to the context of utterance are syntactically triggered).<sup>3</sup>


I think it is clear that any theory which wishes to retain the status of a formal theory of meaning is constrained to give the first answer here: any appeal to the context of utterance must be syntactically triggered. After all, it seems that this is just what it is to pursue a formal account: to think that everything which can be found at the semantic level can be traced to the syntactic level (that we don't get anything 'for free', as it were, at the semantic level). It is this requirement which underpins the problem all moderate formal semanticists have with unarticulated constituents, since these are supposed to be instances of semantically relevant but syntactically unrealized context-sensitivity. In the face of arguments for unarticulated constituents the moderate formal semanticist has two options: either she can argue that, perhaps contra to initial appearances, there really are no such semantically relevant but syntactically unarticulated constituents, or she can argue that, perhaps contra to appearances, every appeal to context really is syntactically triggered.<sup>4</sup> Clearly there is much more to be said on this issue, but since all moderate formal semanticists seem to agree on what the formal stance commits them to here, I want to leave this question to one side and turn to the second, more controversial, aspect of context-dependence.

As noted above, a second way in which we can ask about the kinds of appeal to a context of utterance which are licensed by a moderate formal theory concerns the *kind* of contextual features to which appeal may be made. Again, it seems to me that we have two possible answers here: on the one hand, the relevant features may be entirely non-perspectival, or objective, features of the context, like the identity of the speaker, and the time or the place of utterance. On the other hand, the relevant features may be perspectival or intentional features, such as the mental states of the speaker. Now, objections may be raised concerning my choice of terminology here – for it might be objected that the mental states of a speaker are every bit as much an objective feature of reality as things like times and places. However, even if one doesn't like the terminology, it seems that the distinction (between the mental states of agents and the

<sup>3</sup> King and Stanley (forthcoming) call the first kind of appeal to context 'weak pragmatic effects' and the latter 'strong pragmatic effects'.

<sup>4</sup> For an advocate of the former route see Borg (forthcoming), while for an advocate of the latter see Stanley and Szabo 2000.





other aspects of a context of utterance) is clear enough.<sup>5</sup> Thus our question concerns which of these two types of contextual feature can be semantically admissible from within moderate formal semantics – should the contextual parameters in a moderate formal semantics like Kaplan's be limited to objective features of the context, or can they include perspectival, intentional features as well? Well, it seems that many moderate formal theorists think that *both* kinds of feature are admissible.<sup>6</sup> However, it seems to me that a moderate formal semantics is in fact constrained to reject the second option here – speaker intentions cannot be appealed to from within a moderate formal theory.<sup>7</sup>



There are two reasons for this: first, admitting intentional features of the context of utterance seems to me to run counter to the overarching aims of formal semantic theorising (roughly, to capture the repeatable, code-like, normative aspects of linguistic meaning) and to run the risk of collapsing the difference between formal accounts and use-based approaches. This sort of worry seems particularly pressing given the kind of moderate formal semantics adopted by Stanley and Szabo 2000, whereby every common noun turns out to be context-sensitive. Thus to determine the semantic content of any sentence-type relativised to a context of utterance, if that sentence contains a common noun then it turns out that we need to look to the intentions of the speaker uttering the sentence prior to arriving at anything truth-evaluable. Yet this seems pretty close to at least some ways of characterizing a use-based approach to meaning. At the very least, it seems to me that a moderate formal theory which makes no appeal to the intentional states of the speaker in specifying semantic content counts as a stronger form of formal theory than one which takes speaker intentions as pervasive at the semantic level. So it should provide the starting place for the moderate formal semanticist – a position to move from only should it prove impossible to maintain.


However, the main reason for rejecting appeals to speaker intentions from within a formal semantic theory concerns the clash between the deductive, computational nature of a formal theory (e.g. delivering canonical derivations of truth-conditions) and the abductive nature of inferences

<sup>5</sup> Indeed drawing such a distinction is not novel: Kaplan discussed the question of whether or not speaker intentions were relevant to determining the referent of a demonstrative (see Kaplan 1977: 524–26; Kaplan 1989: 582–84), and the distinction has been deployed by a number of other theorists, including DeRose 1992: 916 (who labels it 'objective vs. subjective') and Bach 1999: 72 (who labels it 'wide vs. narrow').

<sup>6</sup> See, amongst many others, Kaplan 1989; Stanley and Szabo 2000; King 2001.

<sup>7</sup> Bach 1999, and forthcoming, agrees with this claim, although he also denies a key assumption I want to preserve, namely that semantics is in the business of delivering truth-conditions, rather than, say, truth-conditional fragments.










concerning speaker intentions. Speaker intentions are potentially unrepeatable, rich, nebulous features of the context and as such are formally intractable – we simply lack any properly formal way to model grasp of speaker intentions. This point comes into sharp relief when we consider a potentially key motivation for adopting some kind of formal semantics, namely its compatibility with a (Fodorian) modular model of the mind.

Although this is a familiar point, let us just recap on what a modularity theory of mind claims. According to the modularity hypothesis, the apparently unitary body of the human mind is actually composed of some number of discrete, relatively autonomous processors or modules. Opinions differ as to the precise nature of these modules, but on a Fodorian account they consist of encapsulated bodies of information together with rules operating only over that information (Fodor 1983). The processes modules deploy are syntactically-driven, deductive, computational processes – they look not to the content of a representation but only to its syntactic properties. This contrasts with the kinds of processes which Fodor claims take place within the agent's non-modular 'general intelligence'. These processes are abductive (i.e. inference to the best explanation), non-demonstrative reasoning processes, looking not just to the syntactic structure of representations but also to their content and context. Now, we might ask ourselves, if an account of meaning makes appeal to the intentional states of the speaker (whether this appeal is syntactically triggered or not), what does this entail for the kinds of processes involved in coming to grasp that meaning? I think that the answer to this question is as follows: any theory which admits appeal to speaker intentions as relevant in determining literal meaning involves the kind of abductive, non-demonstrative reasoning process which Fodor places beyond the reach of genuine modules. Thus any theory of semantic content which appeals to speaker intentions cannot form part of a modular language faculty.

To see this, consider the way in which we go about offering an intentional explanation of any kind of action. So, imagine that you see Jack putting money into a drink-dispensing machine. Then you might reason as follows: 'Jack is putting money in that drink-dispensing machine. The best explanation for this action is that Jack is thirsty and wants a drink; therefore Jack is thirsty and wants a drink.' Clearly this is a non-demonstrative piece of reasoning and it is susceptible to the influence of an open-ended range of contextual factors. For instance, say that Jack has just finished playing squash with Jill, who is now sitting on the floor, with her head in her hands, looking very red. Then the best explanation of Jack's action might well be that Jill is thirsty and Jack is getting a drink for her. Alternatively, if Jack has said that he needs change, the best explanation for his action might be that he thinks the machine will give him change.





So a range of contextual factors can matter in determining an intentional explanation and there is no determinate boundary at the outset on which facts could turn out to be relevant (and this is just to note that the process looks distinctly unencapsulated). As Recanati notes:



A distinguishing characteristic of pragmatic interpretation is its defeasibility. The best explanation we can offer for an action given the available evidence may be revised in the light of new evidence. Even if an *excellent* explanation is available, it can always be overridden if enough new evidence is adduced to account for the subject's behaviour. It follows that any piece of evidence may turn out to be relevant for the interpretation of an action. In other words, there is no limit to the amount of contextual information that can affect pragmatic interpretation. (Recanati 2004: 54)


On the other hand, however, although it is true that there is no limit in principle to the range of evidence that can affect pragmatic interpretation, it is also clear that, in practice, only a very small subset of all the things an agent knows are actually brought to bear in any instance of such problem solving. As Fodor puts it:

[I]t's just got to be possible to determine, with reasonable accuracy, the impact of adopting a new belief on one's prior commitments without having to survey those commitments in their totality. *Whole theories* can't be the units [of] computation any more than they can be the units of confirmation, or of assertion, or of semantic evaluation. The totality of one's epistemic commitments is *vastly* too large a space to have to search if all one's trying to do is figure out whether, since there are clouds, it would be wise to carry an umbrella. (Fodor 2000: 31)

Or, to put it another way, the totality of one's epistemic commitments is vastly too large a space to search if all one is trying to figure out is whether Jill's utterance of 'I'm too tired' meant that she is too tired *to go to the cinema tonight* or that by her utterance of 'That's mine' she intended to refer to the racquet Jack was holding. The global yet selective nature of the processes underpinning intentional explanation (their potential sensitivity to anything one believes paired with their limited search space in practice) means they cannot be treated as simple computational processes, since they must be sensitive to more than the syntactic properties of representations; neither local syntax (i.e. of the specific representation itself), nor global syntax (i.e. of the entire belief system), underpins these kinds of processes.

So, if we think that delivering truth-conditions for the kind of sentences given above involves an appeal to speaker intentions, it just has to be the






case that the processes by which that content is arrived at are abductive.<sup>8</sup> That is to say, a semantic theory which admits speaker intentions as relevant in determining truth-conditional content could not form part of a formal, Fodorian language faculty. Now, clearly, for this point to lend genuine support to moderate formal semantics (as I've construed it), we would need to see that Fodorian modularity provides the best model of our cognitive architecture with respect to linguistic comprehension.<sup>9</sup> I think there is indeed good evidence for this claim (evidence concerning the speed of linguistic comprehension, its domain specificity, the encapsulation of semantic content, and the patterns of breakdown it undergoes), but this lies way beyond our current endeavour. For all that is necessary to recognize now is that semantic level appeals to speaker intentions take us outside the formal domain. Intentional interpretation is a non-modular, abductive process and as such is outside the domain of a genuinely formal, computational theory of meaning.<sup>10</sup>

In conclusion: if the problems of context do lead us to reject strong formal semantics and move instead towards some more moderate kind of approach, then we face the question of the kinds of contextual appeal a formal theory can make. I suggested that this issue arises in two distinct way: first, there is the question of *when* an appeal to context can be made

<sup>8</sup> One possible response here would be to claim that the kinds of speaker intentions most moderate formal semanticists want to take as semantically relevant (e.g. the referential intentions of a speaker using a demonstrative) are somehow more primitive than full-blown speaker intentions and thus may be accessed in some more brute kind of way, akin to perception. However, this claim looks tendentious – though it is clearly obvious that, say, a speaker is pointing in a given direction, this won't be sufficient for determining reference. Pointing in manner  $p^*$ , though something which is perceptually available, doesn't tell us if the agent intends to refer to that dog, that colour, that collar, or that area, etc. Making these kind of fine-grained distinctions seems to rely on precisely the type of intentional explanation appealed to above.

<sup>9</sup> A discussion of massive modularity would also obviously be relevant here, since advocates of this kind of approach characterize modules in a much more relaxed fashion than Fodor 1983. Two points on this, however: first, it's not at all clear that modules of any kind are the right kinds of things to solve the problem of abduction, and since it is this problem which is at issue here, it is also not clear that the move to massive modularity will be of any help at this juncture. Secondly, even if we were to adopt the massive modularity picture it would remain a live issue as to whether any cognitive structures met the more stringent conditions for being a Fodorian module, and the evidence might point towards linguistic understanding falling into this category.

<sup>10</sup> This in itself does not contravene McDowell's (1980: 138) point that 'semantics, formal or otherwise, is secondary to the theory of communication in at least this sense: if there were no subject matter for the theory of communication, there would be nothing for semantic theories to be about'.




and, second, there is the question of the *kind* of contextual feature to which appeal may be made. I've tried to argue here that it is characteristic of genuinely formal theories to give minimal answers to both these questions. First, moderate formal semantics demands that all appeals to a context of utterance must be syntactically triggered. Secondly, however (and more controversially) I've argued that a moderate formal semantics can only admit semantic-level appeals to objective features of a context of utterance. For an appeal to intentional features takes us beyond the computational, syntax-driven processes proper to formal semantics and into the realm of abductive, inference-to-the-best-explanation reasoning. Now, of course, it is a further question whether or not a semantic theory with such minimal contextual pretensions can still be satisfactory (e.g. deliver truth-conditional content), but all I've argued for here is that the search for such a minimal theory is well-motivated, since speaker-intentions lie outside the reach of formal semantics.

*The University of Reading*  
*Reading RG6 6AA, UK*  
*e.g.n.borg@rdg.ac.uk*

### References

- Borg, E. Forthcoming. *Minimal Semantics*. Oxford: Oxford University Press.
- Bach, K. 1999. The semantics-pragmatics distinction: what it is and why it matters. In *The Semantics/Pragmatics Interface from Different Points of View*, ed. K. Turner, 65–83. Oxford: Elsevier Science.
- Bach, K. Forthcoming. Context *ex machina*. In *Semantics vs. Pragmatics*, ed. Z. Szabo. Oxford: Oxford University Press.
- DeRose, K. 1992. Contextualism and knowledge attributions. *Philosophy and Phenomenological Research* 52: 913–29.
- Fodor, J. 1983. *Modularity of Mind*. Cambridge, Mass.: MIT Press.
- Fodor, J. 2000. *The Mind Doesn't Work That Way*. Cambridge, Mass.: MIT Press.
- Kaplan, D. 1977. Demonstratives. In *Themes from Kaplan*, ed. J. Almog, J. Perry and H. Wettstein, 481–564. Oxford: Oxford University Press.
- Kaplan, D. 1989. Afterthoughts. In *Themes From Kaplan*, ed. J. Almog, J. Perry and H. Wettstein, 565–614. Oxford: Oxford University Press.
- Katz, J. 1977. *Propositional Structure and Illocutionary Force*. New York: Thomas Y. Crowell.
- King, J. 2001. *Complex Demonstratives: A Quantificational Account*. Cambridge, Mass.: MIT Press.
- King, J. and J. Stanley. Forthcoming. Semantics, pragmatics, and the role of semantic content. In *Semantics vs. Pragmatics*, ed. Z. Szabo. Oxford: Oxford University Press.
- McDowell, J. 1980. Meaning, communication and knowledge. In *Philosophical Subjects: Essays Presented to P.F. Strawson*, ed. Z. Van Straaten, 117–39. Oxford: Clarendon Press.



- Recanati, F. 2002. Unarticulated constituents. *Linguistics and Philosophy* 25: 299–345.  
 Recanati, F. 2004. *Literal Meaning*. Cambridge: Cambridge University Press.  
 Stanley, J. and Z. Szabo. 2000. On quantifier domain restriction. *Mind and Language* 15: 219–61.

## *A solution for Russellians to a puzzle about belief*

SEAN CRAWFORD

1. According to Russellians, there is one proposition expressed (in a context) by the two sentences ‘George Eliot is a novelist’ and ‘Mary Ann Evans is a novelist’, given that George Eliot is Mary Ann Evans (Salmon 1986; Braun 1998). It has, as its constituents, the woman George Eliot and the property of being a novelist, and can be represented as an ordered pair:

<George Eliot, being a novelist>.

Since belief is a binary relation holding between a believer and a proposition (ignoring time and tense), there is likewise one proposition expressed (in a context) by the two sentences ‘Ralph believes George Eliot is a novelist’ and ‘Ralph believes Mary Ann Evans is a novelist’. It too can be represented as an ordered pair:

<<Ralph, <George Eliot, being a novelist>>, believes>.

though, when it comes to focusing on the content of Ralph’s belief, which is the present topic, it is perhaps more perspicuous to represent it as

Believes (Ralph, <George Eliot, being a novelist>).

On this view, the two sentences:

- (1) Ralph believes George Eliot is a novelist
- (2) Ralph believes Mary Ann Evans is a novelist

cannot diverge in truth value, since they express the same proposition. The problem for the Russellian (or the ‘Millian’) is that a puzzle of David Kaplan’s (1968–69: §11) seems to show that they can diverge in truth-value and that therefore, since the Russellian holds that (1) and (2) express the same proposition, the Russellian view is contradictory.

Imagine that Ralph has read *Middlemarch* and *The Mill on the Floss* by one ‘George Eliot’ but is too benighted to realize that the author is

