EMMA BORG

SAYING WHAT YOU MEAN: UNARTICULATED CONSTITUENTS AND COMMUNICATION

In this paper I want to explore the arguments for so-called 'unarticulated constituents' (UCs). Unarticulated constituents are supposed to be propositional elements, not presented in the surface form of a sentence, nor explicitly represented at the level of its logical form, yet which must be interpreted in order to grasp the (proper) meaning of that sentence or expression. Thus, for example, we might think that a sentence like 'It is raining' must contain a UC picking out the place at which the speaker of the sentence asserts it to be raining. In §1 I will explore the nature of UCs a little further, and, in §2, suggest that we can recognise two different forms of argument for them in the literature. I will argue that ultimately neither is convincing, and they will be rejected in §3 and §4 respectively. The claim will be that, though the need for an appeal to such things as time and speaker are undoubtedly necessary in order to specify what a speaker said in a given context, advocates of the semantic relevance of UCs have failed to hold apart crucially different aspects of our understanding: first, the difference between knowledge of truth-conditions and the knowledge that truth-conditions are satisfied; second, the difference between knowledge of meaning and the understanding of communicative acts. Instead of ceding contextual information the kind of semantic role envisaged by advocates of UCs, we should, I will argue, see it as part of a theory of speech acts.[1]

I will suggest that what we need to recognise here is the proper division of cognitive labour, for once this division is in place we can recognise the role and function of the information attributed to UCs, and its crucial relevance to communication, without ceding it semantic value. Sketching a model of our cognitive architecture which can underpin this stance, and showing why it might be thought independently attractive, will be the task of §5. Clearly, then, although the main focus of this paper rests with UCs, there are some big issues hovering in the wings here, and perhaps before we turn our attention squarely on the main target it would be in order to say why I think discussion of UCs cannot be had in isolation from these bigger issues.

The reason, as I see it, is that arguments for UCs are part and parcel of a particular perspective on semantic theorising, one which is over-ambitious about the aims of a semantic theory. Consider the tasks we might expect an adequate semantic theory to fulfil: on the one hand, we might be concerned that such a theory explain quite 'low level' linguistic data, such as the meaning possessed by basic lexical items and how, given this base, our language displays properties like systematicity and productivity, which have been made so much of in recent linguistics and cognitive science. On the other hand, however, we might think that an adequate theory should do this and *more*,

R. Elugardo and R. J. Stainton (eds.), Ellipsis and Nonsentential Speech, 237–262.

say incorporating the knowledge required for our general communicative competence, or perhaps even underpinning our epistemic or metaphysical access to the world.[2] The information the theory need contain to achieve the first, limited function might, it seems, be given by a recursive, truth-conditional theory of the kind initiated by Davidson, where the input to the theory is given by structural descriptions of (for the most part) the surface level constituents of sentences;[3] but it is pretty clear even from the outset that such a theory will not take us very far in satisfying the latter kind of constraint. What then makes the existence of UCs seem so compelling to so many theorists, I will argue, is a certain conception of the role of a semantic theory: if we approach a semantic theory from an over-ambitious perspective, then, regardless of the force of any particular argument for UCs, their existence will come to seem inevitable. While, if we limit our ambitions, the semanticist can and should do without such additional, covert elements which receive no linguistic representation. It is for this reason, then, that once I have argued against the specific arguments for UCs (§§1–4), I will go on to say something briefly (§5) about the role and function proper to a semantic theory, arguing that such a theory should be of the limited form which makes the need for UCs otiose. However, let us begin on more solid terrain by examining the nature of UCs and the arguments on offer for their existence.

(1) WHAT ARE UNARTICULATED CONSTITUENTS?

Determining the precise nature and role of UCs unfortunately proves a little harder than was suggested in the introduction. For although all theorists in this area seem happy to agree that a feature like the location where it is said to be raining in an utterance of 'It's raining' constitutes a paradigm case of a UC, the precise account which makes this the case can differ. In this section I want to outline two distinct notions of what might constitute an unarticulated constituent: the first stems from Bach, and the second from theorists such as Sperber and Wilson, and Recanati. Having distinguished these positions, we will then focus our attention on the second — for it is in this latter guise that UCs have recently come to prominence and it is under this guise that their existence proves problematic for standard truth-conditional approaches to meaning. In order to state these two approaches clearly, however, it will be useful initially to consider the way in which syntax and semantics in general relate.

At the start of his *Talk About Beliefs*, Crimmins sketches a principle of compositionality he calls 'full articulation'. This is easiest to state if propositions are viewed as structured entities (containing individuals and properties), then full articulation holds that each element of the proposition literally expressed by an indicative utterance of a sentence must itself be the content of some component expression of that sentence.[4] To put matters crudely, the idea is that the constituents of the proposition expressed by the sentence are *exhausted* by the contributions of the component expressions in that sentence and their mode of combination — we don't get anything 'for free' at the propositional level. Now, as Crimmins notes, for this constraint to play a role in practice we need to clarify what counts as a 'component expression'. One thought might be that component expressions are equivalent to vocalised (or orthographic) words, so that propositional constituents can be read (more or less) directly from surface form. However, there are cases which seem to show that this version of the articulation constraint is too strong. For consider cases such as 'syntactic ellipsis', where the proposition a sentence can be thought literally to express contains more constituents than can be

traced to the surface form of the sentence itself (emerging instead from the linguistic context in which the sentence is to be found).[5] So, consider the following (where the material inside brackets is unpronounced):

(1) A: 'Has Bill gone?'
B: 'Yes, he has [$_{VP}$ gone]'
(2) A: 'Whose dog is that?'
B: 'It's Bill's [$_{NP}$ dog]'

In both of these cases, B's utterance appears to express a proposition containing a constituent not found at the vocalised, surface form level. However, because the additional material *is* present in the immediate linguistic environment of the utterance, and can be simply recovered from here, it is often assumed that the unvocalised material can be treated as a genuine constituent of the sentence B produces. The material is present at the syntactic level, it is suggested, but elided at the surface level. If this is possible then such cases do not contravene the principle of full articulation, though they do require it to operate at the level of syntactic, not surface, form.

A second respect in which syntactic component expressions may diverge from straightforward accounts of surface form has recently been explored by Taylor and Recanati.[6] For they suggest that a full description of the syntactic constituents of a sentence should include those elements represented at what we might call the 'sub-syntactic' level. So, for instance, say we have a transitive verb, the lexical entry for which tells us that it possesses 'slots' for two arguments. If only one argument place is filled in the surface form of a particular utterance of that expression, the presence of the other argument place is nevertheless guaranteed by the sub-syntactic form. For instance, take the lexical entry for a verb like 'kicks', treated as a transitive verb with one argument place for the agent and one for the object (so that the form of the relation is '*x* kicks *y*'). Then, if we get a surface level description of a sentence utilising this expression, but with only one argument place explicitly filled (e.g. 'John kicks'), the syntactic level description of that sentence will nevertheless supply the second argument place, with an existentially bound variable acting as a placeholder, yielding 'John kicks something' or '($\exists x$) John kicks x'.[7]

The principle of full articulation is obviously at its strongest if it holds between the surface form and the proposition expressed by the sentence; however, at this level, as we have seen, the constraint seems too strong. For we need to allow elements to appear at the level of proposition expressed which are not mirrored by component expressions at the level of surface form. Perhaps, then, the right place to state the articulation principle is not at the level of surface form; perhaps instead we should require each element in the proposition expressed by an indicative utterance of a sentence to be contributed by a component expression in that sentence's syntactic form. With this principle in mind, it now becomes easier to state our distinct definitions of unarticulated constituents, according to whether they reject or accept full articulation; so let's turn to this task now.[8]

The first take on UCs allows them to figure as elements of a thought entertained on hearing the utterance of a given sentence, but holds them to be quite extraneous to the proposition literally expressed by that sentence. Full articulation, mapping syntactic form to semantic form, is endorsed for the proposition literally expressed by the utterance of a given sentence, though it is explicitly recognised that the thought entertained on hearing such an utterance may have a content which diverges from the proposition

expressed by that sentence.[9] Specifically, the thought entertained may contain constituents not found in the linguistic item under consideration. This is the position we find with respect to Gricean implicature and (more importantly from our respects) in Bach's notion of impliciture (to which we return in §4). It should be clear that this first definition of UCs does not threaten the principle of full articulation nor the project of standard truth-conditional approaches to semantics. It merely highlights the fact that literal, truth-conditional semantic analyses may not be the only kinds of analyses of meaning we are interested in in communicative exchanges. So, if UCs are thought to undermine the standard programme of formal semantics, they must be understood in some other way.

Our second definition of UCs, however, does directly threaten standard truth-conditional approaches to semantics. On this model (familiar from, for instance, Sperber and Wilson's 1986 Relevance Theory) the output of the formal (context-independent) portion of our semantic theory stands in need of several pragmatic refinements prior to arrival at the proposition a sentence-token expresses: for instance, as well as the processes of disambiguation and reference assignment, familiar from the standard truth-conditional picture, the output of the formal theory may also require the introduction of novel pragmatically triggered elements, through a process Sperber and Wilson have termed 'free enrichment'. These elements, these 'unarticulated constituents', are not mirrored by elements in the surface form or the logical form of the sentence under consideration, they simply figure in the proposition expressed, which may be analysed truth-conditionally, to give the 'explicature' (Sperber and Wilson's term) of the sentence uttered — the semantic analysis of what the sentence, as produced in that context, means. This approach denies the principle of full articulation, whether it runs off the surface form *or* the syntactic form of a sentence: *not all* propositional elements are contributed by component expressions at the syntactic or surface level. For Sperber and Wilson, amongst others, pragmatic mechanisms come to figure, not just 'post-semantically' (as in Gricean implicature), but as an inherent part of the truth-conditional analysis; we must engage in pragmatic reasoning *prior to* arriving at something which is truth-evaluable.[10] It is for this reason that Recanati has labelled such positions 'truth-conditional pragmatics' and it should be clear that such approaches are *incompatible* with the standard truth-conditional approach to semantics.

It seems, then, that we have two quite different definitions of 'unarticulated constituents' to hand:

i. a UC is an element which figures in the thought entertained on presentation or production of a sentence, but which is not relevant to the semantic analysis of that sentence (i.e. it does not figure in the proposition expressed by that sentence token).
ii. a UC is an element required for grasping the proposition literally expressed by an indicative utterance of a sentence, S (i.e. S's explicature), yet an element which receives no linguistic (i.e. syntactic) representation.[11]

I want now to set aside the former definition and concentrate instead on (ii). For on this analysis the contextual information captured by UCs comes to figure in the proposition expressed by a sentence. If this is correct, then the standard, formal approach to semantics, which sees the specification of truth-conditional content as a process which is context-independent, must be rejected.[12] In what follows I want to suggest

that, contrary to the arguments of Recanati, Sperber and Wilson, Carston, et al, there is no role in a semantic theory for these additional, pragmatically introduced elements; rather we should see them as part of a broader theory of thought, independent of language. So let us turn now to the argument for the existence of UCs, understood along the lines of (ii) above.

(2) THE ARGUMENTS FOR UNARTICULATED CONSTITUENTS

Since unarticulated constituents (UCs) are, *ex hypothesi* (and on any definition) not 'visible to the naked eye', as it were, arguments for their existence must be of an indirect form: they will be vindicated as theoretical postulates necessary to explain and underpin the recognised behaviour of the sentence. Prima facie, arguments for their existence then fall into two distinct camps:

(1) For at least some sentences, given just the syntactic constituents of the sentence, no truth-evaluable proposition can be recovered (without UCs the sentence simply lacks truth-conditions).

(2) For at least some sentences, given just the syntactic constituents of the sentence, the wrong truth-conditions will be recovered (truth-conditions based solely on verbalised constituents do not fit our intuitions about the circumstances in which the sentence will be true or false).[13]

A clear advocate of the first form of argument for UCs is Carston 1988, who argues that, for a range of cases, we simply lack anything truth-evaluable if we restrict elements of the 'explicature' (the literal content of what is said) to what we find on the surface of the sentence or at the level of logical form.

For instance, considering an utterance of 'She didn't get enough credits and can't continue', in a context where the most plausible interpretation of the utterance is: 'Jane didn't pass enough university course units to qualify for admission to second year study and, as a result, she can't continue with university study. Jane is not feeling at all happy about this', Carston writes that:

> The question then is which aspects of this interpretation are explicitly expressed (that is, part of the explicature) and which are implicit (implicated)? The disambiguation of 'get' and 'units' and the referent assignment of 'she' are surely part of the explicit content, while the assumption that Jane isn't feeling happy is surely implicit. But what about 'to qualify for admission to second year study' and 'with university study', which enrich and complete the two clauses of the conjunction, and the 'as a result' linking the two conjuncts. Are these part of what is explicated or what is implicated? Since they are not given linguistically, one might think they must be implicated, but then what is the explicature of the utterance? It must be 'Jane didn't pass enough university course units and Jane cannot continue (something??)'. It's not clear that this constitutes a propositional form, that is, it isn't possible to specify what conditions in the world must obtain for it to be true.[14]

It seems clear that Carston here envisages an argument along the lines of (1): without the addition of relevant contextual information, despite the fact that it does not figure in the syntactic form, the sentence is not truth-evaluable, we cannot specify the conditions under which it would be true. Furthermore, even some theorists who would disagree with Carston about the existence of truth-conditionally relevant UCs, agree

with the claim under consideration here, viz. that at least some indicative sentences are semantically underdeterminate. For instance Bach writes:

> An (indicative) sentence is semantically underdeterminate if it fails to express a complete proposition — determine a definite truth-condition — even after ambiguity and vagueness are resolved and indexical references (including the time of the utterance) are fixed . . . .In these cases what the conventional meaning of the sentence determines is only a fragment of a proposition or what I call a *proposition radical*; a complete proposition would be expressed only if the sentence were elaborated somehow, so as to produce a *completion* of the proposition.[15]

Again, then, the thought seems to be that, unless we are willing to take the sentence as possessing more content than it superficially appears to have, we will simply be unable to assign a truth-value. The sentence as it stands is simply not truth-evaluable.

However, as noted above, this is not the only kind of argument possible for UCs, for it may be that, concentrating just on the explicitly represented elements of the sentence, we get something which *is* truth-evaluable, but that what we get is, in some sense, the *wrong* truth-conditions for what is said.[16] Arguments of the form of (2) are most evident in the discussions surrounding quantifier restriction. So, for instance, consider the following exchanges:

(3) A: How was the party?
B: Everyone was sick.[17]

(4) A: I've invited my boss for dinner.
B: But there is nothing to eat!

(5) A: Can I let Fido in from the garden?
B: Yes, the door is closed.

In (3) it seems B's utterance is true just in case everyone at the party was sick (as opposed to, say, everyone in the world); in (4) B clearly doesn't mean an unrestricted claim concerning the lack of food, but something like 'there is nothing to eat in the house' or 'there is nothing appropriate and available to eat'. While in (5), the special case of quantifier restriction that arises with respect to definite descriptions treated as quantified phrases, it seems B's utterance may be true in a situation where the door to the street is closed, even if one or more internal doors are open. If we take our T-sentences to be given simply by the overt elements of the sentence, we must treat (B) in each case as saying something (trivially) false, whereas our intuition in each case is that they have spoken truly.

Another set of cases which seem to lend support to the second argument for UCs can be found in examples like that made famous by Partee:

(6) A: I turned off the oven.

Here, unless there is some implicit reference to a time (and on the assumption that the speaker has turned off the oven more than once in the past) the speaker seems to be saying something trivially true; but this seems wrong. Certainly, the natural way to interpret (6) is along the lines of:

(7) I turned the oven off then.

The problem in (3–6) is not, then, that the sentences uttered entirely lack truth-conditions, but rather that they lack *suitable* truth-conditions. This disparity between our intuitive judgements of what is said (i.e. the conditions under which what is said will be true) and the paucity of the verbalised content of the sentence is again thought to provide evidence for the existence of unrepresented but semantically relevant constituents.[18]

In what follows I want to reject both these forms of argument for the existence of syntactically unrepresented but semantically relevant UCs.[19] Contrary to the first argument given above, I will argue (in §3) that even sentences like 'She can't continue' are truth-evaluable, though we need to hold apart the truth-conditions a sentence possesses and the actual situation which serves to make it true on any given occasion (which in turn may link to judgements of what is pragmatically communicated in that situation). While contrary to the second argument above, I will suggest that we have good reason to take the notion of *appropriateness* as a non-semantic one — one which goes hand in hand with determining what a speaker can or could convey in a given context, but not what a sentence literally means. Although I do not want to query our intuitions about what is said by the speaker in these cases, I will argue that judgements about what is said are of little help in determining what the sentence literally expresses. (Thus I will be *rejecting* the assumption, common from Grice on, that there is some privileged notion of 'what is said' which is informative as to the precise semantic content of the original sentence; seeing why this is so will be one of the tasks of §5.) So, let us turn now to the first argument for UCs.

(3) (SOME) SENTENCES ARE NOT TRUTH-EVALUABLE WITHOUT UCs

The initial argument, endorsed in Carston 1988, claims that some sentences are not truth-evaluable without appeal to UCs. Prima facie, however, it seems that the opponent of UCs — e.g. someone who advocates a disquotational T-theory running (more or less) simply off the surface constituents of sentences — might wonder what exactly the problem is supposed to be here. Why, she might wonder, can't we simply disquote the sentences in question to yield theorems of the form:[20]

a. 'It is raining' is true (in L) iff it is raining.
b. 'Jane can't continue' is true (in L) iff Jane can't continue.[21]

Now, what is clear with truth-conditions of this form is that they don't specify a unique set of conditions which must pertain in order for the sentence to be true; or, better, they allow a range of more specific conditions each of which would serve to make the sentence true.[22] For instance, in the case of (b), we might envisage a range of possible situations, each unilaterally an instance of Jane's failure to continue; e.g. a world where Jane can't continue sleeping, a world where she can't continue running, and a world where she can't continue university education, to name but three.

However, this permissiveness in the precise conditions which make the object language sentence true doesn't immediately seem either problematic or particularly unusual. Take the sentence 'Jane is happy', which, we might think, is a less likely candidate for containing UCs than either (a) or (b). Given a disquotational T-theory we arrive at something of the form:

c. ‘Jane is happy’ is true (in L) iff Jane is happy.

Yet (c), no less than (b), fails to uniquely constrain the range of possible situations in which the object language sentence will be true. A world in which Jane is happy because it is her birthday but not because her boyfriend has left her, or where she is happy now but not five minutes ago, or where she has never been unhappy, are all worlds which serve to make the object language sentence true. And verifying whether or not the sentence is in fact true will involve finding out if any one (or more) of these possible situations is actual; i.e. it will require determining the precise conditions which, in this instance, satisfy the truth-condition. Yet, I would suggest, this is *no* different to what happens in a case like (b): determining the truth or falsity of the object-language sentence (b) will require finding out which (if any) of a range of possible situations are actual. Finding out whether Jane is happy, then, involves undertaking exactly the same kind of investigation as finding out whether Jane can’t continue, it is just that we might think (speaking somewhat crudely) that there is a ‘broader’ range of situations which would make it true that Jane can’t continue than there are which satisfy ‘Jane is happy’.

Yet we clearly need much further argument to show that this intuitive difference in range must result in a difference in meaning, i.e. that there is some recognisable degree of variation in the possible situations which serve to satisfy a truth-condition, below which no introduced unarticulated constituents are needed, but above which the semantic requirement for UCs comes into play. Who, we might wonder, is responsible for setting this line, what exactly are the parameters of difference which it is supposed to be measuring, and what do we do with borderline, disputed or vague cases? To the extent that this proposal can actually be (non-metaphorically) understood, it seems entirely arbitrary and artificial. Of course, one option here would be for the advocate of UCs to deny an initial assumption we made above, viz. that ‘Jane is happy’ is not a good candidate for containing UCs. Perhaps ‘Jane is happy’ is precisely on a par with ‘It is raining’ or ‘Jane can’t continue’, requiring UCs to specify the location, duration and kind of happiness Jane is enjoying. However, any intuitive support for the imposition of UCs seems to dissolve when we turn to sentences like ‘Jane is happy’; although a multitude of different situations (perhaps an infinite number) can satisfy the T-sentence:

‘Jane is happy’ is true iff Jane is happy

this does not, I suggest, in any way encourage us to enrich the semantic content of the sentence in order to narrow down this number of situations. Yet as for ‘Jane is happy’ so for ‘Jane can’t continue’, *unless* the advocate of UCs can convince us that there is some principled distinction between the two.[23] So, it seems, either the advocate of UCs pursuing this first line of argument (given as (1) in §2) must be willing to draw a line at some point, below which UCs are not required (however, in this case they face serious questions concerning how to make this border appear non-arbitrary), or they find themselves on a slippery slope which can only end with the requirement that the literal meaning of every sentence be exactly as precise as the particular worldly conditions used to verify it on a given occasion of utterance. Yet neither of these positions seems appealing.

From the start, then, it seems to me that the burden of proof rests with the advocate of UCs to show us what is wrong with the kind of liberal truth-conditions a formal theory would supply, for though they don’t tie the world down to a unique state of

affairs nor, it seems on closer inspection, do many other sentences.[24] Yet there is no intuitive appeal to the idea that all these other sentences (like 'Jane is happy') are good candidates for UCs. In response to this kind of argument, however, I think the advocate of UCs can marshal further putative problems with the kind of truth-conditions currently under consideration, problems which may still show us that liberal truth-conditions are unacceptable. The three objections I envisage here are as follows: first, there seems to be a problem regarding the fact that the world may both satisfy and fail to satisfy a given truth-condition at the same point in time; second, we may worry about how assessments of truth and falsity actually get made for sentences; third, it may be objected that the contextual conditions in play simply demand semantic accommodation through the role they are playing. I want to explore each of these putative objections in turn, but the conclusion will be that none of them support the claim that at least some sentences are non-truth evaluable without UCs.

A first objection to liberal truth-conditions concerns the recognition, given voice to by Perry, that at any given time, bits of the world may satisfy a truth-condition like (1), while other bits don't. Of his son's utterance of 'It's raining', Perry writes: "What my son said was true, because it was raining in Palo Alto. There were all sorts of places where it wasn't raining".[25] Of course, it is unarguable that the fact that it is raining in Palo Alto serves to make Perry's son's statement true in this situation. However, the claim we might envisage being made here is that this fact requires semantic recognition because otherwise we will be faced with a liberal truth-condition which one area of the world satisfies while another does not. The question we face, then, if we adopt truth-conditions like (a) or (b) is: is a world where it is raining in Palo Alto but not in London a world where the sentence 'It is raining' (without UCs) is true or false?[26]

Recall, however, that the argument currently under consideration is whether or not the sentence is *truth-evaluable* without appeal to UCs (not yet whether the truth-conditions are appropriate), and nothing in the recognition that a part of the world may satisfy the condition in question, while another part of the world does not, serves to show that the sentence is not truth-evaluable: a world where it is raining *anywhere* is, I would suggest, a world where the sentence 'It is raining' is true. For the sentence 'It is raining' to be false, it has to be the case that there is no (current) instance of raining going on at all.[27] If the speaker wanted to assert something which further constrained the set of circumstances which would make her sentence true, then she could and should have done this; but knowing how communication proceeds, she did not feel it was necessary in this case. (I'll return to the question of how much interlocutors can assume in communication, without explicitly asserting, in §5.)

Furthermore, that the sentence 'It's raining' must at least sometimes be analysed along the lines of (a) is reinforced by consideration of cases like the following: say we are concerned to measure the level of rainfall worldwide, perhaps in the light of fears about global warming.[28] To do this we set up a machine which rings a bell whenever there is an instance of rain anywhere in the world. Hearing the bell, it seems I may utter 'It's raining', aiming to express just the proposition that it is raining in some, quite unspecified location. There is no precise location where I wish to assert the rain is falling, nor does the recognition that it is not raining in very many places seem to affect the truth of what I say. So, the recognition that different parts of the world may (concurrently) satisfy or fail to satisfy a given truth-condition is not as yet reason to reject liberal truth-conditions like (a) and (b). However, there is a related problem which begins to surface now, for if truth-conditions really do (for the most part) run

off the explicit contents of our sentences then we will be left with a vast range of truth-conditions whose actual satisfaction we cannot verify. Is the sentence 'it's raining', as uttered by S at t, literally true or false, we might ask? Well, without a relativisation to a place, it can turn out to be extremely hard to tell.

Once again, however, it's not clear that the advocate of standard truth-conditional semantics should be unduly worried by this claim. For it seems that the claim made by this kind of approach to semantics is that grasp of meaning is grasp of truth-conditions, not knowledge of whether those truth-conditions are satisfied, nor possession of a method by which to discover if those truth-conditions are satisfied; to think otherwise is, I believe, to fall prey to a kind of creeping Verificationism. What we are allowing is that the competent interlocutor can grasp the truth-conditions of the sentence, she knows how the world would have to be for the sentence to be true. To think that, in addition to this, the agent must be in a position to ascertain whether or not that condition is satisfied in order to count as understanding the meaning of the sentence is to run together notions of meaning and verification which (the history of Verificationist approaches to meaning tells us) are best kept apart. What matters for understanding a sentence is that it have a truth-condition, i.e. that it be (in principle) truth-evaluable, and that the interlocutor grasp that truth-condition, and this is not at all the same thing as requiring that, at any given time, we must be in a position to actually determine the sentence's truth-value. (It should also be born in mind that any such failures to verify concern *only* the proposition literally expressed by a sentence. It is perfectly possible that speakers of these sentences will convey some more precise proposition through their utterance of the sentence in question, and that the truth-value of this pragmatically conveyed proposition will be easily verifiable by interlocutors; a point returned to in §4.)

We need to hold apart knowing the truth-conditions of a sentence (a semantic matter) and knowing whether or not those truth-conditions are satisfied on some particular occasion of utterance (a non-semantic matter). What is obviously the case, given our limited cognitive resources and the speed of communicative exchanges, is that we simply don't have the time or ability to check all possible situations satisfying the conditions on any given occasion; but we should also note that very often we don't have to. Take the sentence 'John went for a walk', which can be made true by a world in which he went for a quick walk by a lake half an hour ago, or by a world in which he went for a slow walk over a bridge two weeks ago (and countless many other worlds as well). To find out if this sentence is true, I will begin by investigating those circumstances which are most likely to have provided the evidence for my interlocutors production of the sentence. If I discover, amongst these relevant alternatives, a situation which makes the sentence true, then I can simply stop there; if my interlocutor is speaking truly, then I can usually expect to find a confirming situation fairly quickly, say discovering that John did indeed go for a quick walk by a lake a short time ago.[29] Clearly, then, attempts to *verify* whether or not a given truth-condition is satisfied may well be something of a limited or curtailed endeavour. Specifically, we may confine ourselves to what seem to be the *relevant* possibilities here. We may decide that, even though a world in which Jane can't continue sleeping is a world in which 'Jane can't continue' is true, it is not a very relevant circumstance for us to investigate. Rather, from the conversational exchange in which the sentence is embedded, it seems that we can figure out a much more relevant set of circumstances to devote our attention to; namely, whether or not Jane can continue university education. But, to repeat, unless we think that meaning is

to be located in our methods of verification, there is simply no argument from the role of a particular place, speaker, type of footwear, etc, in our verification of a sentence's truth-value to the necessary inclusion of such elements in a specification of semantic content.

It seems, then, that the advocate of disquotational T-theories can maintain that disquotation (based solely on syntactic constituents) is adequate for generating semantic content, whilst admitting that the conditions interlocutors look to to verify whether a given truth-condition is satisfied or not are severely curtailed, i.e. that the kinds of features Carston et al want to add to the semantics actually figure outside the semantics in the realm of how agents go about verifying the truth or falsity of a given utterance. In this case, sentences like 'Jane can't continue' are perfectly truth-evaluable (all we need to do is to grasp the appropriate disquotational T-sentence), though verifying the truth-value of the sentence, i.e. determining if its truth-condition is satisfied, may well advert to the kind of contextual information appealed to by advocates of UCs. So, I want to suggest that simple, disquotational T-sentences, like (a) and (b), are perfectly acceptable: first, though they do not pin the world down to a unique state of affairs, we have no reason to expect or require them to do so. Secondly, though this entails that the propositions literally expressed by many sentence tokens will not be verifiable, this only constitutes a problem if we lose sight of the fact that knowledge of meaning is knowledge of truth-conditions, not possession of a method of verification for those truth-conditions.

Finally, however, this brings us to the third and last objection to permissive truth-conditions that I want to consider. For we may worry that, if it's really the case that judgements about the truth or falsity of 'It's raining' stand or fall with how the weather is *here*, or *at X*, as indicated above, then this should be a fact which is reflected in our semantics. Not because without appeal to a place the sentence is non-truth-evaluable, but because without such appeal the sentence is not *appropriately* truth-evaluable. Truth-conditions like (a) and (b), though not ill-formed, are not suitable; they fail to capture our intuitive judgements about when the sentence should be taken to be true or false. The real worry here, then, does not seem to be, as initially suggested, that we simply can't get anything remotely truth-evaluable without appeal to UCs, rather it is to suggest that we can't get anything approaching *appropriate* truth-conditions without appeal to UCs. The objection is that, without the proposed presence of UCs, the only kinds of truth-conditions we can deliver for sentences like 'Jane can't continue', 'It's raining', or 'There's nothing to eat', are ones which fail to capture our judgements about when sentences like these are true or false. This, of course, is the second argument given above for the existence of UCs, so let us turn to this alternative form of argument now.

(4) UCs ARE NEEDED FOR APPROPRIATE TRUTH-CONDITIONS

It seems, then, that the first argument (that some sentences are literally not truth-evaluable) ultimately collapses into the second argument; viz. that truth-conditions based solely on the syntactic contents of (at least some) sentences are in some way inappropriate. What a speaker says when they utter the sentence 'Everyone was sick' is that every person *in some relevant group* was sick, and what the speaker who says 'It's raining' means is that it is raining *in some particular place*. To treat the sentences as

possessing the more general truth-conditions delivered by their overt constituents is to fail to capture what the speaker means, and to make predictions about the circumstances in which the sentence will be true or false which do not fit with our intuitive view of the subject matter. For instance, it is to hold 'Everyone was sick' is false in a situation in which everyone at the party was sick, but where some irrelevant individual, whom no one was talking about or thinking about, was well. Or that 'It is raining' is true when it is bone dry for hundreds of miles around the interlocutors, but, unbeknownst to all participants in the conversation, it is raining in a small corner of Timbuktu. This result, the advocates of UCs object, is unacceptable.

Initially, then, the claim seems to be that someone who utters, say, 'It's raining', can or must be viewed as meaning that it is raining *in X*, etc. However, as we all know thanks to Grice, because a speaker means a proposition, p, by her utterance of a sentence, s, this does not necessarily mean that the sentence uttered should be treated as having the semantic value that p.[30] The speaker who says, ironically, 'It's a nice day' when it's raining, means *it's a nasty day*, but this isn't the literal meaning of the sentence uttered. While the speaker who says pointedly 'someone hasn't handed in their essay again' may mean, and may be taken to mean, that the recalcitrant Jones has failed to turn in work once again, but this isn't what she literally expresses. So, the advocate of a restricted (non-UC) view of semantics might wonder why the cases to hand are any different. Why should we think that, because it is often uncontentious to say that the speaker uttering 'It's raining' means that *it's raining here* (or wherever) that this more informative proposition must give the literal, semantic content of the sentence produced?

I think there are probably three factors at play in the advocate of UCs assumption that the richer proposition must give the literal meaning of the sentence. First, the kinds of cases which the advocates of UCs appeal to intuitively look pretty different to typical cases of Gricean speaker-meaning. In the latter, we have an intention on the part of the speaker to say something non-literal (they intend to be ironic, hyperbolic, metaphorical, etc), whereas in the kinds of 'underdetermined' utterances focused on for UCs any such non-literal intent is absent. Why, then, think speakers are knowingly uttering literal falsehoods or trivial truths on these occasions, even when they are in possession of a range of perfectly simple sentences which would convey the substantive thoughts they really wish to communicate (i.e. why don't they say 'It's raining here' as opposed to just 'It's raining')? This thought seems especially pressing since our intuitions in these cases tell us that the speakers are in fact doing fine — producing fitting utterances and (often) asserting truths. So, even if the Gricean distinction is right in certain cases, still, the advocates of UCs contend, there is no reason to think these cases are (and every reason to think they are not) instances of speaker-meaning rather than semantic-meaning.

This connects to a second reason to treat the richer propositions as giving the literal meaning (as opposed to being Gricean implicatures), for it often seems both natural and correct to report a speaker who produces a (putatively) underdeterminate utterance using a 'completed' content sentence. For instance, the speaker who says 'It's raining' will usually be reported as having said that *it's raining here* (or wherever). Finally, as we saw at the close of the last section, it seems that when we look at the conditions appealed to in order to make judgements of truth or falsity for sentences as uttered on a given occasion, they are the states of affairs picked out by the richer propositions involving time, place and speaker, etc. If I want to find out whether Perry's son's utterance of 'It's

raining' was true or false I need to consider not how the weather is with me now but how the weather was with him then (even though finding out the latter state of affairs is a much harder task than the former). Yet if this is right, then it seems undeniable that the statement Perry's son made must have contained elements picking up on his particular context of utterance (or the intended context for his utterance), i.e. that the semantic content of the sentence uttered contained UCs.

Let's take these three points in order: first, the claim that these cases are radically different to paradigm Gricean examples of speaker-meaning. It's obvious that speakers producing such sentences lack the kind of explicit non-literal intentions Grice appealed to. Yet this is not necessarily to concede that the speakers in these cases are intending to be taken literally; that is to say, we shouldn't think of non-literality as necessarily exhausted by cases where the speaker is trying to be ironic, metaphorical, etc. For instance, as Bach has pointed out (in discussion of what he terms 'implic*i*ture'):

> [T]here are many sentences which are almost always used non-literally as elliptical for other sentences. For example, "Ed doesn't look tired, he is tired" would likely be used with a suppressed "merely" before "look" to be inferred by the hearer, since the speaker would not be stating that Ed does not look tired but is tired anyway. Similarly, if I say "I drink only Scotch", I would not be stating that I drink nothing but Scotch but merely that the only liquor I drink is Scotch . . . .The phenomenon of elliptical speech is commonplace; indeed, it often seems stilted not to suppress words that can easily be inferred as expressing part of what one means, as opposed to what the uttered sentence means.[31]

The claim that, in utterances like 'It's raining', the speaker is not trying to be *explicitly* non-literal (in the sense of trying to be ironic, etc) only entails that the speaker is not being non-literal if this is the only kind of non-literality we allow. What Bach points out in the above quote (and elsewhere) is that this last claim is extremely tendentious: we seem to allow a wide range of cases of non-literality, stretching much wider than the mere intention to speak metaphorically or ironically, etc.[32] It seems instead that a speaker can be viewed as speaking non-literally just in case there is a divergence between the thought the speaker wants to express and the literal meaning of the sentence produced, and this kind of phenomenon happens frequently, especially when there are elements already in play in the context which it would be stilted to repeat in one's speech, even though they form a part of the content of the thought to be conveyed. So, is there any evidence that speakers in the kinds of cases under consideration here are being non-literal in this (broad) sense?

Clearly the answer to this question is 'yes', for there is evidence that interlocutors are willing to hold apart literal sentence-meaning and speaker-meaning even in these cases. For instance, in an utterance of "I will go to the store" it always seems open to the mischievous speaker, on being chided to actually go, to reply that she did not say *when* she would go and that she merely meant to express the proposition that at some time in the future she would be visiting the store. No doubt such a speaker contravenes all sorts of communicative or conversational constraints, but it doesn't seem that she explicitly *contradicts* herself (as must be the case if her original sentence literally meant that, for some specific value of *t*, she would go to the store at *t*).[33] Or again, take the cynical response to 'Everyone is coming to my party' of 'Oh really? Will the Queen be there?' — the respondent here may be charged with being pedantic and uncharitable, but surely not with failing to understand the literal meaning of the English sentence. The retreat to the general proposition acquired from the surface contents of the sentence may be pedantic, and a speaker who insists on such unhelpful interpretations will quickly

prove an exasperating interlocutor, but the mere fact that we allow such retreats, without charging the speaker with inconsistency or failure to grasp the meaning of the sentence, seems to demonstrate that we have here precisely the kind of sensitivity to speaker-meaning versus sentence-meaning outlined by Grice. So, the first argument against treating the richer propositions as non-literal can, it seems, be deflated.

Moving to the second point, concerning indirect speech reports, it seems that the shared intuition — that, for instance, an utterance of 'It's raining' can often be reported using the richer content sentence 'A said that it's raining *at l'* — can be accommodated without embracing semantically relevant UCs. For it seems that this intuition may be best viewed as concerning not the literal meaning of linguistic expressions in natural language, but speaker's exploitation of these signs in successful communication, as highlighted in reported speech. That is to say, just because, in many contexts, it is entirely natural to report a speaker utilising a content sentence which is richer than the sentence originally uttered, this does not mean we need refine our semantic evaluation of the original sentence to incorporate every element present in an acceptable content sentence of an indirect speech report. For, in general, it seems quite clear that facts about reported speech cannot be used in any straightforward way to demarcate facts about semantic content.[34] The first thing to notice is that the move from proposition expressed to correct indirect speech reports is not one:one. A single utterance can always be adequately reported by a number of indirect speech acts. So, 'It's raining' as uttered by S, at time t and location l, may be reported in (at least) the following ways:

S said that it is raining.
S said that it is raining where she is.
S said that it is raining at l.
S said that it was raining at l on t.

Furthermore, given the right context of reporting, the utterance may support a range of more 'liberal' indirect speech reports, like:

S said that it was raining 50 miles south of the Grand Canyon.
S said that it was nice weather for ducks at t in l.
S said that the drought was over.

So, if we were to assume that facts about indirect speech limn facts about semantic content, we would have to allow that a single sentence possesses an indefinite number of distinct semantic contents, depending on the range of acceptable ways in which it may be reported.[35] Yet, with concerns surrounding the systematicity and creativity of natural language (and the constraints of language learning) in mind, this seems totally unacceptable. It simply seems wrong to think that part of the semantic content of 'It's raining' could include reference to droughts or ducks.

Of course, the natural move for the advocate of UCs here is to claim that it is not *every* indirect speech report which is informative as to semantic content, but only some subset of them (e.g. those which capture what the speaker 'really' said, in some sense).[36] However, as Cappelen and Lepore have stressed, once one starts reflecting on how permissive indirect speech reporting can be, the idea that it can tell us anything useful about semantic content becomes extremely doubtful. For instance, given the right

context of utterance and report, the content sentences in indirect speech reports may swap co-referring or synonymous terms from the original sentence (e.g. exchanging 'John' for 'that boy'), and they may exchange referring terms for quantified noun phrases ('John' for 'the oldest boy in class').[37] We may also allow the omission of conjuncts or disjuncts ('p & q' reported by 'S said that p'), and the picking up of implicatures ('someone hasn't done the washing-up' reported by 'John complained that Jill hadn't washed up again', or Blair's claim that 'I will endeavour by the office of this government to bring once again within the direct control of the Nation those systems of public transportation that form the lifeline of so much of this country's wealth and well-being' reported as 'Blair said that he wants to renationalise the railways'). Yet given this degree of liberality, it seems very hard to see how the discrete subset of indirect speech acts which are intended to be genuinely informative as to semantic value are to be distinguished.[38] Rather, it seems, facts about reported speech *per se* entail very little about what meaning should be ascribed to the original sentence uttered. Thus, there is no direct move from the intuition that 'It is raining at l' may be a correct report of an utterance of 'It is raining' to the theoretical claim that the former gives the correct semantic analysis of the latter.

To recap: it seems that the first and second motivations for assigning UCs semantic relevance in order to arrive at appropriate truth-conditions can be dissipated. For, on the one hand, there is evidence that the cases in question do fit the speaker-meaning/sentence-meaning divide introduced by Grice (since we are willing to take the speaker's rejection of assigned, richer propositions — like *it is raining here* — and their retreat to the more general proposition yielded by syntactic constituents alone — e.g. *it is raining* — as non-contradictory and legitimate, though almost certainly conversationally improper and pedantic). While, on the other hand, it seems that the data here properly resides with facts about reported speech (viz. the unarguable fact that speakers can be correctly reported using content sentences which overtly appeal to such elements as speaker, location and time); but, as I have tried to show, for this undisputed fact to be relevant here, we need to assume an extremely close connection between how a speaker can be reported (i.e. what the speaker succeeded in communicating) and the literal meaning of the sentence uttered, a connection which in general does not seem to hold. Though how a speaker can be reported must be *in some sense* constrained by the sentence she produces, the assumption that we can extract facts about semantic meaning from facts about reported speech seems wrong. In §5 I will offer an explanation of why this is the case, sketching a view of the cognitive architecture of the agent which makes it clear why we cannot hope to begin with facts about speaker-meaning and hope to move from there to facts about sentence-meaning. However, the claim for now is simply that an advocate of the semantic relevance of UCs owes us much further argument from the claim that contextual information figures in indirect speech reports to the idea that such information figures semantically in the initial sentence produced.

However, the advocate of UCs is not to be silenced yet, for as noted at the start of this section, there is a third argument she may appeal to in rejecting the pragmatic explanation of these cases. For perhaps the motivation for ceding the richer propositions semantic relevance can be found in consideration of the conditions under which we seem willing to judge the sentences in question true or false. As noted above, it seems that the speaker saying 'It's raining', in a context where the relevant location is X, will be judged to have spoken truly *if it's raining at X* and falsely even if it's raining

elsewhere (and not at X). While the speaker who says 'There is nothing to eat' may be judged to have spoken truly, despite the absence of global famine. So, how can we claim contextual information is irrelevant to sentence-meaning when the conditions under which utterances of these sentences are held true or false are just the kind of constrained conditions delivered by the incorporation of UCs? It is not only that the kind of information appealed to by UCs figures in correct indirect speech reports, but also that it figures in our assessments of the truth and falsity of the original utterance; what more evidence could we need, the advocate of UCs will object, to grant the information a semantic role?

I think there are two points to notice in respect of this argument: first, we need to bear in mind the distinction between knowledge of truth-conditions and the verification of those truth-conditions, and, secondly, we need to ask ourselves *which* proposition interlocutors will be most interested in verifying the truth of in any given context — will it be the literal, semantic content expressed, or will it be the proposition (or propositions) the speaker wants to (and succeeds in) communicating? On the first point: as we saw in the last section, we need to be very clear that the conditions interlocutors appeal to to verify a sentence are not necessarily identical to the truth-conditions of the sentence produced. I may verify the truth of 'John went for a walk' by finding out that he went for a slow walk over a bridge, or a fast walk beside a lake, but neither the speed of the walk nor the route taken (need) figure as part of the semantic content of the sentence. Similarly, I may verify an utterance of 'It's raining' by seeing that there is a downpour outside my window, or being told by a reliable source that there is a light drizzle in Palo Alto, but, I contend, this gives us no reason to think the weight of waterfall or the place where it is falling figures in the literal meaning of the sentence produced. To think otherwise would be to demand the literal meaning of the sentence produced be precisely as fine-grained as the particular condition used to verify it; but we have no reason to think every sentence we produce must specify a completely unique way the world must be in all its myriad detail. Furthermore, such a position would run roughshod over any principle of 'semantic innocence' we might have, by seeing the contents of sentences like 'John went for a walk' as containing concepts like bridges and lakes, strolls and wanderings.[39] Rather it seems that verifying the truth of a sentence is simply not the same thing as understanding the truth-conditions of that sentence, and it seems that the final argument for UCs is guilty of running together these two notions: maintaining that just because we appeal to a specific condition in determining the truth-value of the sentence, this condition must be part of the semantic content of the sentence produced.

Furthermore, even though I believe we should hold apart the notions of truth-conditions and the verification of truth-value, it still seems that the opponent of semantically relevant UCs can accommodate the crucial role contextual features play in understanding communicative acts. To see this, we need to be clear about the nature of the debate here, for matters are somewhat delicate. The issue here is not 'does contextual information have a role to play or not?' (a question to which all parties will answer in the affirmative), but 'does this information have a peculiarly semantic role to play?' (the issue is one of division of labour). The opponent of UCs can grant relevant contextual information a crucial role to play in understanding and verifying the truth of *what is said* (non-semantic) by the utterance of the sentence in the given context, even whilst denying it a semantic role in the literal meaning of the sentence produced. That is to say, they can explain why we may judge 'It's raining'

as false when it is not raining *here*, though it is raining (at some irrelevant) *there*; or why we judge 'There is nothing to eat' true, even when there is a well-stocked food shop nearby, for in these cases we are judging not what is literally expressed but what is communicated. Contextual information is of crucial importance for understanding what speakers mean, but this is not to say that it must have an inherent role in the literal meaning of those sentences speakers use to communicate what they mean.[40] What advocates of semantically relevant though syntactically omitted elements should recognise, I think, is that there is a perfectly standard discrepancy between sentence-meaning and speaker-meaning, and that features vital for determining both the truth-conditions and the truth-value of the latter need not be in any way relevant for determining the former.

So I believe we should reject all three of the proposed reasons for treating the richer propositions resulting from the inclusion of contextual material as semantically relevant, treating them instead as quite standard cases of speaker meaning (as opposed to sentence meaning). Yet if this is right, then the second argument for UCs fails: *appropriate* truth-evaluation is a pragmatic matter, thus the fact that UCs are required for this goes no way towards establishing their semantic relevance. Combined with the failure of the first argument for UCs (namely, that some sentences are non-truth-evaluable without UCs, rejected in §3), it seems that we are left with no compelling argument for the existence of semantically relevant, though syntactically unmarked, constituents. Finally, however, this brings us back to the bigger issues touched on at the outset; for it seems to me that to reject the existence of such unarticulated constituents it is not enough simply to reject the specific arguments for them. For it seems that, in actual fact, for many theorists the real motivation for UCs comes from the embracing of a particular perspective on semantic theorising — a perspective which makes such elements almost inevitable. The suggestion we have to consider now is that, regardless of any specific argument for the existence of UCs, if we want a semantic theory which is in any way adequate, we will simply be forced to accept the existence of UCs; arguments about the semantic relevance of UCs are ineliminably connected to arguments about the role of a semantic theory itself.

The advocate of UCs apparently sees the semantic realm as primary — as responsible not only for our understanding of linguistic items, but also for our understanding of what speakers can use these items to say. The assumption seems to be that a semantic theory which is not sensitive to the range of thoughts conveyed in a communicative exchange must fail as a theory of meaning. Alternatively, opponents of UCs see the semantic theory as contributing just *one* element to the understanding of what was said by the speaker, with elements such as knowledge of the speaker, knowledge of the context, identification of the referent, etc, forming equally crucial, though non-semantic, elements. Thus so-called 'pragmatic' features (a label which, I think, really serves just as a 'catch-all' term for non-semantic knowledge) are different to, but absolutely no less vital than, semantic features. In the final section, then, I want to consider briefly wider questions concerning the semantic/non-semantic divide, examining how we might construe the boundary between language and thought, how it might be crossed, and in which direction. The argument will be that we have no reason to amalgamate all the information required to understand communicative acts as properly part of the semantic theory and that, with the more constrained view of semantics in place, the need for semantically relevant though syntactically unarticulated constituents drops away.

(5) THE BOUNDARY BETWEEN LANGUAGE AND THOUGHT

The positive view I want to put forward is that understanding a language is just one, necessary but far from sufficient, step on the road to understanding linguistic communicative acts. Understanding a language is not, as Wittgenstein told us, understanding a way of life; rather it is understanding a constructed code, a system of representation with finite basic parts and recursive rules, which can be used by speakers to express (elements of) their thoughts. Of course, this claim as it stands is probably uncontroversial: all theorists, advocates of UCs, truth-conditional pragmatics and standard formal semantic theories alike, want to recognise the important role of postsemantic, pragmatic features in affecting speaker's meaning. However, where the view to be advocated diverges, I think, from the commonly accepted claim, is in the degree of responsibility attributed to, and the range of, non-semantic information. The thought is that semantic interpretation yields only an extremely minimal level of understanding and that what we need to do to build up to anything like an adequate understanding of a communicative act is to subject this semantic interpretation to a vast range of further information we possess concerning the world and one another.[41] We need, to put matters somewhat hyperbolically, to move from language to thought.

To make matters more concrete, let's borrow from the picture made familiar by theorists like Chomsky and Fodor. Within this framework, then, what I want to claim is that agents possess a language faculty containing discrete bodies of information, say, orthographics/phonetics, syntax, and semantics. The semantic information contained in the language faculty is, however, of a quite minimal kind, namely just what is required to explain the kind of low-level semantic facts given in the introduction (meanings of primitives, properties like productivity, etc). On its own, then, the language faculty is not equipped to explain fully our communicative competence. To know what someone has said (non-semantic) by an utterance of a given sentence, an interlocutor needs to begin with the calculation of the literal meaning of the sentence produced, but this information is then fed out of the language faculty and into what we might call an agent's 'generalised intelligence' (in current jargon, sometimes the 'central processing unit'). It is at this point that the specifically semantic information becomes subject to a vast range of other kinds of information possessed by the agent, including the output of the perceptual system, commonsense psychology, commonsense physics, etc. The point, which is often paid lip-service but not, I think, always fully appreciated, is that what matters in (even linguistic) communication is *just as much* what is not said as what is.[42] To arrive at an understanding of what is said a great deal of language-independent information must come together; though we *start* with an understanding of the meanings of words and their modes of combination, we proceed almost automatically to an assessment of what that literal meaning itself means in the current context and in the mouth of the current speaker. The literal meaning may be enriched, altered, rejected or refined in the light of an agent's non-semantic knowledge. It is for this reason that trying to recover the purely semantic contribution from an assessment of what is communicated is so difficult: we cannot start with the coalescence of all these different features and hope to drag out the semantic contribution from here, for it has been submerged within our general understanding. Though there is a function which takes us, in any particular case, from the sentence produced and the context in which it is uttered, given our background grasp of the world and one another, to the proposition communicated, we

can do no more than offer a functional definition of this operation in terms of its input and output.[43]

To understand what is literally meant by 'It's raining' all we need to know is the meaning of the parts of the sentence (as indicated by the surface elements of the sentence) and their mode of combination, however to know what is communicated by an utterance of this sentence we need to know so much more. For a start, we need to understand some crucial facts about language-based communication, such as that when a speaker comes to conveying a particular idea in a given context she may choose to use words and phrases which do not entirely match her thought. This may be because the thought is (for her) inexpressible — she simply cannot find the right words for it; or it may be because she wishes to flout some conversational rule to a given end — perhaps she wishes to be ironic or metaphorical; or it may be because at least some of the information she wishes to convey is already in the public domain, as it were, so that she can use a short-hand linguistic version of what she means to communicate. It is this 'conversational short-hand' that I want to suggest is in play with the sort of contextual information appealed to by UCs.[44] In addition, given that her grasp of communicative practices tells the interlocutor that information from the wider context of utterance, or background information she possesses about objects or people, will be relevant for determining what thought the speaker means to convey, she needs to know *which* non-semantic features of the context and her background knowledge are relevant for determining what is said, i.e. that in an utterance of 'It's raining' it is more useful to determine which location the speaker has in mind than the kind of rainfall she believes to be occurring, though sometimes determining factors like the speaker's attitude to the rainfall may be equally important. Furthermore, we need to know non-semantic information about the concepts deployed in the literal meaning itself, such as that two drops of water probably don't constitute an instance of rain, or that if it is raining in an area of drought then the drought is over;[45] and we need to know non-semantic information about how agents usually act in response to rain, e.g. that if it is raining then any picnics will be cancelled, or that in the rain people tend to use umbrellas.

The point I want to stress is that the semantic contribution to judgements of what is said forms just one (crucial) part of a much bigger picture, and that without the bigger picture semantic meaning is an impoverished thing. To know a language, if one doesn't know about the world or one another, is not yet to know very much.[46] Thus to position contextual elements as necessary to understanding communicative acts, though not necessary to understanding semantic content, is not to undervalue these additional elements, rather it is just to recognise that they are playing a different (though equally important) role in coming to understand what a speaker means by her utterance of a given sentence. This difference in role may be obscured by the fact that it is both natural and immediate to move from understanding the sentence to understanding what the speaker of the sentence conveys in a given context, together with the fact that, once we have arrived at this latter meaning, it is extremely difficult to retrace our steps to discover the purely semantic contribution to the communicative act (as natural language speakers we are adept at crossing the boundary from language to thought and back again, but as theorists the interaction of different aspects of our knowledge remains a very poorly understood domain).

Yet this is not to say that the distinct contributions should not be held apart: we know the kind of information a semantic theory must contain, I suggest, because we

know the kind of data a semantic theory has to explain, viz. the systematic nature of linguistic understanding and our ability to produce an indefinite range of sentences despite our limited cognitive resources. Furthermore, we have some idea of how this data might be explained, i.e. by positing a recursive, truth-conditional theory as responsible for semantic understanding. Thus it is from *this* perspective that we can isolate the semantic contribution of sentences to judgements of what is said (non-semantic); the mistake made by (some) advocates of UCs is to try to *begin* with judgements of what is said and abstract semantic contributions from there. At present, this latter task is simply beyond us: we cannot in any particular case work back from the result of processes utilising our 'generalised intelligence' to discover the specific contribution of the language faculty, because the kind and amount of other information which figures in a calculation of what is said is simply too vast and too complex.[47] To repeat: if we understand the task of a semantic theory (as I think we should) to be explaining features like the productivity and systematicity of natural language, and how an infant can come to acquire such a language, then there is no place for syntactically unrecognised but semantically relevant UCs. In general, the input to such a theory will be simply a structural description of the surface level features of the sentence. This leaves a great deal still to be said about how we understand the communicative acts in which linguistic items may play a part, for such an approach holds out no hope of a semantic theory coming to serve as a general account of communicative competence. That is to say, such an approach still leaves the door wide open for 'unarticulated constituents' understood in the first way given in §1 (stemming from Bach), for this is just to recognise that the thoughts properly engendered by an utterance of a sentence, S, may diverge from the semantic content of (i.e. proposition expressed by) S to a greater or lesser extent.

In conclusion, then, I have argued that we should reject both the specific arguments for, and the wider perspective which underpins, the move to embrace truth-conditionally relevant but syntactically unarticulated constituents. In the first place, we should reject the claim that (many) sentences are not truth-evaluable without the appeal to such constituents: for by holding apart specification of truth-conditions from verification of truth-value, we can see that this claim is unfounded (the argument of §3). Secondly, we should recognise that the richer propositions containing additional contextual information are well treated as forms of *impliciture*: first, since interlocutors do recognise a distinction between literal and speaker meaning, even in these cases, and, second, because facts about reported speech seem to tell only indirectly on facts about semantic content (the arguments of §4). Finally, I have suggested that we have no reason to lump together all the apparently disparate knowledge required for understanding a communicative act under the general heading of 'semantic'. Rather, we do far better to retain a more austere view of the task of a semantic theory, seeing it as required to explain some quite precise features of our linguistic understanding, but allowing that it contributes just one, necessary but far from sufficient, element to our understanding of communicative acts. Yet within this perspective on semantic theories there is simply no need for syntactically unarticulated but semantically relevant constituents. When I work out that by your utterance of 'It's raining' you mean to convey the thought that it is raining *here* what I learn is something about you, using information about you, our speech community and our current (conversational and other) context; what I learn is not something about the meaning of our language.[48]

NOTES

[1] This position is not novel — it is, for instance, advocated in Bach 1994a and by Cappelen and Lepore (manuscript). Furthermore the general strategy should be familiar to all from such arguments as Kripke's rejection of the semantic relevance of referential definite descriptions. However, I hope the precise arguments deployed against the 'semantic relevance' camp are new.

[2] Theorists who endorse the epistemic role for a semantic theory include those who, following Russell, see fit to posit a special class of linguistic items which require 'acquaintance' (or a similar privileged epistemic relation) to understand. While those endorsing the metaphysical role of semantics include those who predict that we can read our ontology in some way from our language.

[3] The interpolation 'for the most part' is needed to provide the degree of latitude required for cases like syntactic ellipsis, to be explored in §1.

[4] Crimmins 1992, pp. 9–10.

[5] See Elugardo and Stainton 2003 for a somewhat more extended discussion of ellipsis, especially with respect to non-sentential cases. For a challenge to the idea that there is syntactically real though unvoiced material even in cases like (1–2) see Dalrymple (this volume).

[6] Taylor 2001; Recanati 2002.

[7] If we allow that sub-syntactic features are relevant to semantic articulation then, as Recanati stresses, many of the cases standardly treated as instances of unarticulated constituents will not be genuinely unarticulated at all (articulated, as they are, sub-syntactically). I will return to this point in the next section.

[8] It is worth noting in passing a third possible explanation for the kind of material typically accorded to UCs, which stems from Perry 1986 (where I believe the term 'unarticulated constituent' first appeared). Though Perry himself opts for something more akin to the second definition to be given below, he begins (1986, p. 147) with a discussion of information which figures, not as a propositional constituent, nor even as an element of the thought entertained by the agent, but instead as parameter against which the proposition a sentence expresses gets assessed for truth or falsity. It might be thought that a similar treatment, where contextual information is introduced by the character or lexical rule associated with a linguistic item (and is thus not directly incorporated into the truth-conditional content of the uttered sentence) could be available for all the kinds of information attributed to UCs. While I do not have the space to discuss this approach here, two points should be noted: first, such a move will result in a *vast* increase in the number and kind of contextual parameters supposedly introduced by lexical items, which may not be credible. Second, however, even if such a move could be made to work, it would not be in conflict with the standard conception of truth-conditional semantics, and thus does not form a genuine opponent in this respect.

[9] Matters are delicate here, for it may be objected that, in general, it is wrong to think of sentence *types* as expressing propositions, rather we must speak of sentence types *relativised to a context* (to account for the resolution of indexicality). The question then is whether we can think of a context in this respect along broadly Kaplanian lines (i.e. as consisting of a set number of contextual parameters — e.g. speaker and time — which can be settled independently of investigation of richer notions like speaker intentions), or whether the notion of context we need is a far fuller one which itself introduces the kind of elements typically appealed to for UCs. As Rob Stainton has urged on me, this is a crucial issue here, but I hope it is clear that I construe 'standard truth-conditional semantics' as committed to something like the former picture; indeed, the arguments of this paper could be construed as aiming to show that the richer notion of context is not necessary for determining linguistic meaning.

[10] Matters here are additionally complex given that Bach *agrees* that sometimes pragmatic reasoning must be entered into to arrive at something which is truth-evaluable, but *disagrees* with Sperber and Wilson et al about the status of this 'completed' proposition. For Sperber and Wilson it gives the literal explicature of the sentence produced, while for Bach it forms a non-literal, pragmatic supplement to the semantically relevant 'propositional radical' (I return to this point briefly in §2).

[11] Recanati's, 2002, p. 316, chosen definition for these elements is slightly different again. He writes: "*In context*, it may be that the unarticulated constituent is 'required'; but then it is required *in virtue of features of the context*, not in virtue of linguistic properties of the expression-type. A constituent is mandatory in the relevant sense only if *in every context* such a constituent has to be provided (precisely because the need for completion is not a contextual matter, but a context-independent property of the expression-type). This, then, is the criterion we must use when testing for (genuine) unarticulatedness: Can we imagine a context in which the same words are used normally, and a truth-evaluable statement is made, yet no such constituent is provided? If we can imagine such a context, then the relevant constituent is indeed

unarticulated; if we cannot, it is articulated, at some level of linguistic analysis." Clearly, then, Recanati's class of UCs will be narrower than many other truth-conditional pragmatists, though he suggests it still includes elements such as the location of rain in 'It's raining'. This narrower definition of UCs impacts on the range of arguments available to Recanati. To anticipate §2: he will allow arguments of the second form for UCs (viz. that without them inappropriate truth-conditions are delivered), but not the first (viz. that without them sentences are non-truth-evaluable). However, since I aim to reject *both* forms of argument, I will not distinguish Recanati's position in what follows.

[12] In a helpful discussion of the issues involved here, Rob Stainton has tried to convince me that I'm overstating the case here — that processes like free-enrichment are not in tension with the core values of truth-conditional semantics. While it is true that truth-conditional pragmatists and what I'm calling 'standard truth-conditional semanticists' go a long way down the same road together, there do remain some pretty radical differences. For advocates of free enrichment (i.e. the existence of UCs) really do claim that pragmatic enrichment is a necessary precursor not merely of determining a *truth-value* for a token utterance but for delivering the *truth-conditional content* of a sentence token, whereas, apart from a quite constrained set of cases (e.g. delivering a referent for an indexical or settling ambiguity), truth-conditional semanticists simply deny that such a process is necessary. Top-down processing, from pragmatics to semantics, is extremely restricted for the advocate of standard truth-conditional semantics, yet it is a ubiquitous part of linguistic understanding for the truth-conditional pragmatist. Thus if UCs really do exist — if free enrichment plays the role the truth-conditional pragmatist envisages — the standard model of truth-conditional semantics must be mistaken.

[13] (1) and (2) correspond to Bach's distinction between *completion* and *expansion* (see Bach 1994a); where completion occurs if "something must be added for the sentence to express a complete and determinate proposition (something capable of being true or false)" (p. 127), and thus corresponds to our first argument for UCs above. While expansion is the process of what Bach calls 'conceptual strengthening', which is not mandatory (unlike completion), and thus corresponds to our second argument above.

[14] Carston 1988, pp. 33–4.

[15] Bach 1994a, pp. 268–269.

[16] Crimmins and Perry 1989 seem to envisage an argument of this form for UCs representing mode-of-presentation-like entities for belief reports.

[17] This example is from Neale 1990 (*Descriptions*, Cambridge, Mass: MIT).

[18] A third set of cases which may prove relevant here are Travis-type examples (see, for example, Travis 1985; 1996), where judgements of the truth of type-identical sentences seem to depend crucially on some kind of contextual sharpening, e.g. 'John's book weighs 8lbs', where we need to know if it's the book John wrote, the book he owns, the book he's carrying, etc.

[19] Clearly, there are two different ways to argue against syntactically unrepresented but semantically relevant elements: one may argue that, contra first impressions, such elements *are syntactically represented* (see Stanley 2000 for an argument to this effect), or one may argue that, contra first impressions, such elements *are not semantically relevant*. This paper pursues the second strategy.

[20] The following simplified truth-sentences would, of course, need relativisation to speakers and times (perhaps in the form of Higginbotham's 1995 'conditionalised' truth-sentences) in order to handle overt indexicality; see fn. 9.

[21] If we embrace, as I'm inclined to, the earlier specification of 'sub-syntactic' material as influencing the syntactic characterisation of the sentence, then the right-hand side of (b) should in fact take notice of the status of 'continue' as a transitive verb, yielding:

(b*) 'Jane can't continue___ ' is true (in L) iff Jane can't continue something.

The advocate of UCs here would still claim that contextual features have a role to play in determining the correct truth-conditions for a token of 'Jane can't continue', for these must specify *what* she can't continue; but the current line of argument, that this sentence is not truth-evaluable as it stands, would no longer hold. That is to say, to borrow Bach's terminology (see fn. 13), the argument for UCs would now be one concerning *expansion* and not *completion*. Whether or not (a) should be rendered in this fuller form (i.e. whether the lexical entry for 'rain' is akin to 'raineth' or more akin to 'rain somewhere') is a moot point, to be settled by empirical study of our language and its lexicon.

[22] It is sometimes suggested that, if a sentence allows a range of more specific conditions, then it itself must be providing only a truth-conditional schema, or propositional radical. Yet this claim seems far too strong, for no (contingent) proposition graspable by the human mind could be *maximally* specific about the world, thus every proposition will allow *some* more specific conditions to be provided.

[23] It has been suggested to me that an advocate of UCs might respond in the following way: there *is* a relevant difference between 'Jane can't continue' and 'Jane is happy' (assuming for argument's sake that the latter does not require UCs); namely, though both sentences express propositions which do not uniquely constrain the situations which satisfy them, only the former sentence requires contextual supplementation to get a proposition in the first place. But such a response simply seems to me to beg the question: to assert that the difference is that, in the former case you don't, but in the latter case you do, get a proposition, seems to me little more of a statement of faith. We were looking for an *argument* as to why a putative truth-condition like (a) was unacceptable, while (c) was not, and the fact that (a) is not fully determinate about the precise condition which may satisfy it was adduced as a reason. Now however we have seen that this fact holds just as much for (c). If it turns out that very few, if any, sentences actually express fully determinate propositions (in the sense of *fully* describing how the world must be to satisfy them), then the fact that sentences like 'It's raining' don't do this provides *no* evidence that contextual supplementation is required to take us from sentence to proposition.

[24] Here we have a major point of disagreement with Bach, who claims that, for many sentences, the propositions they literally express are incomplete and thus non-truth-evaluable; they are propositional radicals rather than complete propositions. However, at this juncture, I fail to see exactly what the argument is supposed to be. Considering the sentence 'Steel isn't strong enough', Bach asserts 'Notice that [this sentence] does not express the weak proposition that steel isn't strong enough for something or other' (1994b, p. 127), but the argument behind this intuition is unclear.

[25] Perry 1986, p. 138.

[26] A similar thought might seem to be behind Taylor's comment on the same sentence: 'The semantic incompleteness is manifest to us as a felt inability to evaluate the truth value of an utterance of ["It's raining"] in the absence of a contextually provided location (or range of locations)' (2001, p. 53).

[27] Similarly, for the sentence 'It is not raining' to be true, there must be no instance of rainfall anywhere in the world at the current time (so it is not the case that 'It is raining' and 'It is not raining' may both be true at the same time). Advocates of UCs might think it is dishonest to smuggle in reference to the time in these examples, however, I would suggest that this is admissible given the present tense of the sentences.

[28] Recanati 2002 cites this kind of example as well.

[29] Notice, however, that despite the confirming situation containing elements like *being by a lake* or *ten minutes ago*, there is, I would suggest, no temptation to see the literal meaning of the sentence produced as making implicit appeal to these further elements.

[30] See Grice 1967; also Sperber and Wilson 1986.

[31] Bach 1981, p. 238. See also the discussion of 'S-(sentence)-non-literality' in Bach 1994b.

[32] Furthermore, we should note that this point is already endorsed by anyone who accepts a pragmatic, speaker-meaning analysis of so-called 'referential' definite descriptions, where it is rarely the case that the speaker will have explicitly non-literal intentions.

[33] Returning to Perry's case, imagine that you are confused about the current conversational setting and think your son wants to say something about Palo Alto, when in fact he is continuing a conversation about how things stand with your other son in LA. Here the intended proposition is not that it is raining in Palo Alto but that it is raining in LA. Something has gone wrong here, but is it really a semantic failure? Have you failed to understand what your child's words, in that particular concatenation, meant? The suggestion I want to make is that you have not — though there has undoubtedly been a breakdown in communication, this isn't due to your inability to understand English, i.e. to grasp the literal meaning of the sentence, but because you've failed in some other respect, such as keeping up-to-date on which are the relevant objects in the current context.

[34] Cf. Bertolet 1990, and Cappelen and Lepore 1997.

[35] We should also note that facts about admissible reportings vary not only with the context of utterance of the original sentence, but also with the current audience to which the report is being made. For instance, a report of 'Blue is my favourite colour' with 'A said that that is her favourite colour' is clearly acceptable only in a quite specific range of contexts, viz. those that have a sample of blue available for demonstrative indication.

[36] For responses of this kind, though not in the context of the UC debate, see Richard 1998 and Reimer 1998.

[37] Bertolet 1990 allows exchange of co-referring terms but not co-extensive predicates, however little argument over and above the intuition that exchange is permitted in the former but not the latter case, is offered.

[38] See Cappelen and Lepore's 1998 reply to Richard and Reimer. The argument surrounding the semantic irrelevance of judgments of 'what is said' is explored at greater length in Borg 2002.

[39] By 'semantic innocence' I mean some quite general principle requiring the elements we posit within a semantic analysis for a given sentence to be constrained by the syntactic elements we can find in that sentence.

[40] To return to the distinction made in §1, the claim is that we should endorse the first definition of UCs, allowing that there may be elements required to grasp the thought communicated by a speaker which have no immediate counterpart in the proposition literally expressed by the sentence they produce. Some readers have suggested this position tacitly grants all that theorists such as Sperber and Wilson demand, since they are concerned to offer a theory of cognitive content per se (not merely that restricted to linguistic understanding); however, I would disagree (see fn. 41).

[41] In Recanati's 2002 terminology, I wish to defend 'radical literalism'. It is worth being clear, however, exactly where I see the disagreement between my position and that of, say, the Relevance Theorists. For whereas the latter kind of approach seeks some general principle of human cognition, seeing no special boundary between understanding a language and understanding non-linguistic communicative acts, I wish to maintain a clear division, treating linguistic understanding as autonomous from other bodies of knowledge (i.e. pragmatic processes) yet still dealing with fully propositional knowledge (i.e. able to yield determinate truth-conditions for sentences). This fundamental difference in outlook impacts elsewhere, e.g. on questions of the semantic relevance of 'what is said', or analyses of non-literal uses of language.

[42] Notable exceptions here include Bach and Harnish 1979, and Levinson 2000.

[43] See Borg 2000b for further discussion of such putative functional definitions.

[44] It also perhaps yields an initial understanding of non-sentential speech, e.g. holding up an object and saying 'From France' or pointing and saying 'New dress?'. On the current view, though these speech acts would fail to reach the standards of linguistic communication (since the speaker does not produce something which expresses a complete proposition), the speaker might nevertheless 'get her message across' due to the public accessibility of the linguistically elided material. The complete proposition ([This is from France] or [That is a new dress?]) would have the status of a pragmatically conveyed proposition, which interlocutors could recover, despite the linguistic infelicity. For much further discussion of these issues, see Stanley and Szabo 2000, who adopt a treatment sympathetic to the suggestion here, or Elugardo and Stainton 2003 who argue for an alternative view.

[45] Some theorists might object that *all* this information should be construed as genuinely semantic, since the concepts deployed in semantic theorizing are to be individuated by their inferential role, and thus all such inferential relations should be located within the domain of the semantic theory. However, as is well known, such theories face a serious problem with differentiating the essential, meaning-constituting inferences from the (social or idiosyncratic) inferences people are inclined to draw on the basis of past experience, etc. So, to borrow a favourite example of Fodor's 1998, because I am inclined to infer 'x is dangerous' from 'x is a brown cow', because all brown cows I have run into up till now have been dangerous, this of course doesn't mean that this inferential move should be taken to be constitutive of the meaning of the expression 'brown cow'. The suggestion above, then, is that instead we take *no* such inferences as constitutive of meaning; agents make conceptual connections, but these connections come into play as part of the wider, non-linguistic cognitive architecture of the agent and they hold between independently individuated concepts. The meaning of the word 'rain' is just the concept *rain*; what inferences someone is willing to draw on the basis of a deployment of this concept is a function of their experience of the world, not a function of their knowledge of language.

[46] Examination of certain cases of deficit is perhaps relevant here, such as sufferers of Williams syndrome, who have advanced linguistic skills despite being significantly retarded, or those who suffer other impairments of cognitive function which leave language skills untouched. For instance, Pinker 1994, pp. 50–53, describes a case from the psychologist Richard Cromer, where a girl called 'Denyse' talks in detail about the problems with her joint bank account, despite the fact that cognitive disability, as a result of being born with spina bifida, means that she has never had a bank account, cannot read or write, or handle money, and clearly lacks most of the knowledge usually associated with possession of the concept 'joint bank account'.

[47] Recall the quote from Carston 1988, in §2, where she suggested that, given an utterance of 'She didn't get enough units and can't continue', the disambiguation of 'get' and 'units', and the referent assignment to 'she' are "surely part of the explicature", while the assumption that 'she is not happy about this' is "surely part of the implicature". What I want to suggest is that this kind of carving up of semantic and non-semantic understanding is *only* possible if we *don't* start from a characterization of what counts as semantic drawn from considerations about what is required for understanding the communicative act per se.

[48] Thanks are due to Kent Bach, Eros Corazza, Ray Elugardo, Ernie Lepore, Mark Sainsbury, Rob Stainton, and audiences at the Nottingham 'On Referring' conference, and the London Language Reading Group, for helpful comments and discussion. Also to François Recanati for making available draft work on this topic.

REFERENCES

Bach, K., "Referential/Attributive." *Synthese* 49 (1981): 219–44.

Bach, K., "Semantic Slack: what is said and more." In *Foundations of Speech Act Theory: philosophical and linguistic perspectives*, ed. S. Tsohatzidis. London: Routledge. 1994a: 267–291.

Bach, K., "Conversational Impliciture." *Mind and Language* 9 (1994b): 124–162.

Bach, K., and Harnish, R., *Linguistic Communication and Speech Acts* Cambridge, Mass.: MIT. 1979.

Bertolet, R., *What is Said: A theory of indirect speech acts*. London: Kluwer. 1990.

Bezuidenhout, A., "Pragmatics, Semantic Underdetermination and the Referential/Attributive Distinction." *Mind* 106 (1997): 375–410.

Borg, E., "An Expedition Abroad: metaphor, thought and reporting." In *Midwest Studies in Philosophy.* XXV, P. French and H. Wettgtein (eds.). Oxford: Blackwell. 2001: 227–248.

Borg, E., "The Semantic Relevance of What is Said." *Protosociology* 17 (2002): 6–24.

Cappelen, H., and Lepore, E. "On an Alleged Connection Between Indirect Speech and the Theory of Meaning." *Mind and Language* 12 (1997): 278–296.

Cappelen, H. and Lepore, E. "Reply to Richard and Reimer." *Mind and Language* 13 (1998): 617–621.

Cappelen, H., and Lepore, E., "Radical and Moderate Pragmatics: Does meaning determine truth-conditions?" In *Semantics vs. Pragmatics*, ed. Z. Gendler Szabo. Oxford: OUP. forthcoming.

Carston, R., "Implicature, Explicature, and Truth-Theoretic Semantics." In *Mental Representations*, R. Kempson (ed.) Cambridge: Cambridge University Press. 1988: 155–81.

Carston, R., "Explicature and Semantics." UCL Working Papers in Linguistics, 2001.

Crimmins, M., *Talk About Beliefs*. Cambridge, Mass: MIT. 1992.

Crimmins, M., and Perry, J., "The Prince and the Phone-Booth." *Journal of Philosophy* 86 (1989): 685–711.

Dalrymple, M., "Against Reconstruction in Ellipsis." This volume.

Fodor, J., *Concepts: Where Cognitive Science Went Wrong*. Oxford: OUP. 1998.

Elugardo, R., and Stainton, R. "Grasping Objects and Contents." In *The Epistemology of Language*, A. Barber (ed.) Oxford: Blackwell. 2003: 257–302.

Grice, P., "Logic and Conversation (William James Lectures)." In *Syntax and Semantics*, Vol. 3, P. Cole and J. Morgan (ed.) New York: New York Academic Press. 1967: 41–48.

Higginbotham, J. "Tensed Thoughts." *Mind and Language* 10 (1995): 226–249.

Levinson, S., *Presumptive Meanings: The Theory of Generalized Conversational Implicature*. Cambridge, Mass.: MIT. 2000.

Perry, J. "Thought Without Representation." *Proceedings of the Aristotelian Society Supplementary Volume* LX (1986): 263–283.

Pinker, S., *The Language Instinct*. London: Penguin. 1994.

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

Reimer, M., "What is Meant by 'What is Said'? A Reply to Cappelen and Lepore." *Mind and Language* 13 (1998): 598–604.

Richard, M., "Semantic Theory and Indirect Speech." *Mind and Language* 13 (1998): 605–616.

Sperber, D., and Wilson, D., *Relevance: Communication and Cognition*. Oxford: Blackwell. 1986.

Stanley, J., "Context and Logical Form." *Linguistics and Philosophy* 23 (2000): 391–434.
Stanley, J., and Szabo, Z., "On Quantifier Domain Restriction." *Mind and Language* 15 (2000): 219–261
Taylor, K., "Sex, Breakfast, and Descriptus Interruptus." *Synthese* 128 (2001): 45–61.
Travis, C., "On what is strictly speaking true." *Canadian Journal of Philosophy* 15 (1985): 187–229.
Travis, C., "Meaning's role in truth." *Mind* 105 (1996): 451–466.