

# Shared assumptions: semantic minimalism and Relevance Theory

Daniel Wedgwood  
University of Edinburgh

## 1 Overview

Cappelen & Lepore (2005) defend a version of semantic minimalism which they argue to stand in opposition to all forms of contextualism, whether ‘moderate’ or ‘radical’. Relevance Theory (henceforth RT) is included among the latter group and indeed given particular attention as apparently a highly characteristic example of radical contextualism. Cappelen & Lepore (henceforth C&L) produce a number of arguments against both moderate and radical contextualism (many of which have been extensively discussed elsewhere<sup>1</sup>); in this article I am interested in the ways in which they attempt to undermine the radical kind only. Specifically, I question whether C&L’s arguments against radical contextualism do in fact carry any weight when directed against RT—a question with implications for the way they draw their distinction of moderate versus radical contextualism. In particular, I look at C&L’s arguments from *shared content*, one of their primary arguments against radical contextualism, which is expanded in C&L (2006*b*) and applied specifically to RT in C&L (to appear *a*).

My reasons for concentrating on RT, as opposed to the whole of radical contextualism, are the following. First, I find RT to be the most coherent radical contextualist approach; indeed, it may be the only approach that both is truly *radically* contextualist and fully enough worked out to constitute a linguistic framework. This is not the place to justify these claims in detail, but I give some arguments to this effect below (especially in section 5) and otherwise refer the reader to Wedgwood (2005, Chapter 2) and the extensive comparison of RT with other

---

<sup>1</sup>See, for example, Montminy (2006), Atlas (to appear) and the multiple responses in *Mind and Language 21/1* (2006) and *Philosophy and Phenomenological Research* (forthcoming).

approaches in Carston (2002). Second, C&L go to some lengths to distinguish themselves from RT. Third, as I argue below, I in fact see RT and C&L's position as being very closely compatible (despite C&L's claims to the contrary). If this is right, it is a very important result, as it represents the convergence of two major, distinct approaches to natural language meaning, one a philosophical perspective on semantics and one a cognitive perspective on pragmatics, on a rejection of the currently dominant ('moderate contextualist', or 'Gricean') paradigm<sup>2</sup>.

The essence of the C&L's 'shared content' argument is this: if the interpretation of natural language were really very generally context-dependent, as relevance theorists (and any other genuine radical contextualists) claim, then we would be left with no explanation for (a) intuitions about having understood each other and (b) how we regularly share ideas with sufficient precision to allow for all kinds of complex co-ordinated actions between human beings. The only way to explain this ability, they argue, is instead to assume that minimal propositions are conveyed directly and consistently by language, across different contexts.

At the beginning of their (to appear *a*) article, Cappelen & Lepore (henceforth C&L) introduce the notion of shared content as follows:

Speakers share content when they make the same assertion (claim, conjecture, proposal, etc). They also share content when they propose (entertain, discuss, etc.) the same hypothesis, theory, and thought. And again when they evaluate whether what each says (thinks, claims, suggests, etc.) is true, false, interesting, obscene, original or offensive. Content sharing, so understood, is the very foundation of communication.

I identify two key points of confusion in C&L's arguments. The first, dealt with in section 3.1, is the failure to distinguish two kinds of assessment of the content of an utterance: one from the addressee's subjective point of view and the other from the point of view of an omniscient third party (such as a philosopher or linguist). The former plays a key role in RT, while crucial parts of C&L's arguments erroneously appeal to the latter. The second point of confusion concerns different kinds of content. As I discuss in section 4, the phenomena that C&L invoke as evidence against RT involve a kind of content that is not guaranteed to be shared under any theoretical approach. Indeed, their own position on this

---

<sup>2</sup>It is therefore important to note that while in this article I mostly offer criticisms of parts of C&L's work, the underlying motivation for this is to support what I see as the most important aspect of their overall approach.

kind of content—insofar as it is made explicit anywhere—is shown to be equivalent to that of RT in all crucial respects. As such, it is just as well for C&L that their ‘shared content’ arguments against RT fail, as any argument of this kind that succeeded in undermining RT would also undermine their own position.

More generally, I argue that C&L’s whole outlook is in essence very similar to that of RT. This is despite C&L’s conviction that they differ substantially from relevance theorists. For their part, adherents of RT have only rarely engaged directly with C&L. Where they have, differences rather than similarities have been emphasised (Bezuidenhout 2006, Carston 2006). This is understandable enough; indeed, what follows here is also mostly criticism of C&L’s more detailed claims. However, theoretical progress also depends on recognising what approaches share, something that is especially important in this case, given that both RT and C&L stand in opposition to the dominant paradigm in linguistic semantics, which C&L (2005) call *moderate contextualism*. More importantly still, these two approaches are opposed to moderate contextualism in strikingly parallel ways.

Given this, it is particularly unfortunate that Cappelen & Lepore (to appear *a*) misapprehend important aspects of RT. I argue (section 5) that a number of C&L’s criticisms of RT in fact have force only as arguments against moderate contextualism—and that C&L’s arguments against moderate contextualism do not, as they assume, generalise to RT, or to anything that would truly merit the description ‘radical contextualist’. Of course, this is only true as long as RT is genuinely radical contextualist in practice. Below, I discuss a case of moderate contextualist reasoning encroaching on an ostensibly RT analysis and I suggest that parts of C&L’s (2005) discussions can be useful in clarifying key issues and dangers for the relevance theorist.

While I show that C&L’s position in most ways parallels RT closely, there remains one notable difference between the two approaches. This lies in C&L’s notion of *minimal semantic content*, which I argue to be largely equivalent in function to RT’s *encoded meaning* (or, more technically, *underspecified logical form*). However, C&L, unlike relevance theorists, maintain that minimal semantic content is propositional. I discuss this claim in section 4. While this does relate to real differences between C&L and RT, I call into question whether they are always as substantial as they may initially seem—at any rate, they do not affect the arguments surrounding C&L’s criticisms of RT. The key property of minimal semantic content where these arguments are concerned is that it should be conveyed consistently across all contexts. In this respect, RT’s *encoded meaning* does exactly the same job—indeed, somewhat more coherently, as I argue below. Just as importantly, C&L’s view of minimal semantic content leaves them relying on

exactly the kind of processes that RT aims to characterise, while they themselves venture no such characterisation. As such, RT is not only just as well equipped as C&L's approach to deal with 'content-sharing', it is also considerably more useful and explanatory.

## 2 RT and shared content

### 2.1 Inferring and sharing

Sperber & Wilson's (1986) RT is an example of what C&L call radical contextualism, in that it takes the influence of context on propositional meanings to be ubiquitous (an aspect of the theory addressed and defended by Carston 2002 in particular). That is, the intrinsic content of a linguistic form, such as a sentence, is taken to be regularly—perhaps always—sub-propositional. Particular propositional meanings are conveyed only in context, via the addressee's applying certain principles of reasoning about the speaker's intended meaning.

RT is a particularly coherent radical contextualist framework partly because it is not built primarily upon pre-existing semantic traditions (it is after all a *pragmatic* framework). Instead, it is developed from fresh reasoning over some basic observations about the context in which linguistic communication takes place: notably, the fact that such communication takes place between individuals who interact not only with each other's utterances but also simultaneously with their environments—a salient part of which is the other interlocutor and everything that may be known or gleaned regarding that interlocutor.

Another basic observation is that linguistic communication is what Sperber & Wilson term *ostensive communication*: by its very nature, a linguistic utterance conveys the fact that the speaker intends to communicate thereby. This provides the basis for a view of linguistic interpretation as a process of inferring the speaker's communicative intentions. The speaker provides a basis for doing so by formulating an utterance with an awareness of the communicative context. For the speaker this includes information that could reasonably be gleaned about the addressee, which gives rise to expectations about how the addressee will react to certain kinds of linguistic stimulus and certain bits of information. RT therefore recognises communication to be a matter not of simply encoding and decoding messages, but partly of 'mind-reading'—and the goal of RT is to provide an explanation of how this mind-reading is achieved.

RT is thus an attempt to explain how addressees arrive at what, *according to*

*their subjective calculations*, must be a communicator's intended meaning. It is a characterisation of what language users do when faced with any of a huge variety of conceivable circumstances, not of some abstract notion of communication in idealised circumstances. That is, RT is not a semantic theory (though it may impact upon semantic theory rather directly). As such it simply does not deal in objective truth in the way that semantic theories tend to. To criticise RT on the grounds that it cannot guarantee objective identity of content is therefore to miss the point of RT. This in itself constitutes, in broad outline, a response to C&L's 'shared content' argument. To flesh it out requires a brief presentation of key aspects of RT.

## 2.2 The structure of RT and semantic minimalism

In this section, I offer a skeletal outline of how RT works<sup>3</sup>. My aims here are (a) to provide the basis for comparison to C&L's approach, (b) to show how much of RT is based on highly intuitive, minimally controversial observations and (c) to show that RT is not grounded in conventional semantic wisdom but does have ramifications for semantic practice.

RT in effect starts from two basic observations: (i) that what is understood to be communicated by practically any utterance is heavily dependent on context and (ii) that nevertheless people succeed in communicating with each other, commonly without any apparent problem. This means that there must be some set of principles that enables people to recognise the intended meaning with a considerable degree of reliability. I take these observations to be uncontroversial. Certainly C&L would agree with them—indeed, their (2006*b*) article on 'content sharing' is (independently) explicitly structured around observations (i) and (ii).

RT rejects the Gricean picture whereby, in effect, propositions are directly conveyed by sentences uttered (*modulo* somewhat mysterious processes of disambiguation and reference assignment) and the role of inference is confined to the generation of cancellable implicatures. There are two important reasons for this. First, RT recognises the pervasive influence of context on what is understood to be conveyed by a given utterance—hence the *radical contextualist* tag. Second, RT starts out from a much more basic position on what kinds of meaning (or 'content') there must be: simply that some meaning is encoded in linguistic forms and some is inferred. The meaning encoded in a given linguistic expression

---

<sup>3</sup>For proper introductions to the framework, see Sperber & Wilson 1986, Blakemore 1992, Carston 2002, Wilson & Sperber 2004.

is sometimes called its *logical form*, though I prefer to stick to the less loaded *encoded meaning*. Rather than making *a priori* assumptions about the nature of encoded meaning, RT works from the opposite perspective, asking what principles of reasoning enable us to recover a speaker's full intended meaning, given the significant influence of context on what is understood to be conveyed and the fact that people cannot literally mind-read. The question boils down to what people can take to be evidence for communicative intentions.

One very important kind of evidence available to addressees is of course the encoded meaning of the linguistic expression used. But there must also be a means of evaluating the particular significance of this on a given occasion, in the light of some kind of further evidence. Such further evidence is extra-linguistic and may in principle be drawn from any information that is available to the individual. But 'available' here cannot mean 'objectively knowable'—what an individual has to work with is just whatever he has reason to believe, to some extent or other. In RT terms, some set of *assumptions* is *manifest* to an individual, where 'manifest' means roughly 'accessible or deducible on the basis of available evidence'. A crucial fact for communication is that an individual's assumptions will include many assumptions about other people and about the assumptions that are manifest to them. A given pair of individuals 'shares' a given assumption insofar as that assumption is *mutually manifest* to them: i.e. to each person, the assumption is manifest, and so is the fact that it is manifest to both of them. The set of mutually manifest assumptions constitutes those individuals' *mutual cognitive environment*.

All this technical talk adds up to the observation that there are things that we feel we can more or less reasonably assume about the outside world, including about each other and about each other's perceptions of things. This much is surely a necessary basis for any theory of communication. The substance of RT therefore lies not in establishing this view of 'shared content', but in its proposals for how we home in on just those mutually manifest assumptions that are *relevant* to a given act of communication. Clearly, it would be inefficient to the point of paralysing communication to require interlocutors to examine every mutually manifest assumption in the course of producing and interpreting utterances. Equally clearly, one of the problems with this would be that some assumptions that are manifest (i.e. potentially accessible) would require considerable mental effort to actually recover and process.

Reasoning on this basis leads to the fundamental mechanism of RT: interpretation proceeds according to a 'presumption of optimal relevance', where relevance is conceived of as a ratio of communicative rewards to processing effort. These 'rewards' are what RT calls *cognitive effects*: strengthening or weaken-

ing/eliminating existing assumptions or providing new assumptions that interact with existing assumptions to allow for the deduction of further assumptions (known in RT as *contextual implications*)<sup>4</sup>.

Such technical details of RT are less important to the arguments of this article than the general picture that RT presents, as follows. Language users have at their disposal the words and structures of a given language, which clearly carry some very specific kinds of meaning but also may regularly be taken to convey a variety of different things, depending on the assumptions that are brought to bear on the process of interpretation. That is, linguistically encoded meaning constrains but does not determine what is conveyed on a given occasion. Which assumptions are brought to bear on interpretation depends on the relative accessibility of different assumptions and more generally on the processing effort involved in reaching a given interpretation—with addressees only considering more costly lines of enquiry if this is necessary to reach a worthwhile interpretation. Communicators, meanwhile, have to put in some effort to anticipate the outcome of such a strategy of interpretation and to tailor their utterances accordingly (this is not altruistic—the communicator buys his own form of reward thereby: getting his message across).

What does this tell us about shared, or shareable, ‘content’ in RT? As the above discussion makes clear, there are ultimately just two key levels at which one could identify ‘content’ in RT. Fundamentally, RT is a single set of principles for calculating intended meaning given the evidence of (i) the encoded meaning of linguistic expressions used and (ii) the ‘broad context’ (i.e. the mutual cognitive environment of the interlocutors)<sup>5</sup>. Therefore, the only truly distinct notions of ‘content’ that RT makes available are the input to and outputs of this process: encoded and inferred meanings<sup>6</sup>.

This makes RT strikingly similar to C&L’s professed position, which they call *semantic minimalism* combined with *speech act pluralism*. They too recognise just two significant levels of meaning: what they call minimal semantic content

---

<sup>4</sup>The apparently rather complicated definition of contextual implications merely reflects the intuitive fact that new assumptions with no connection to existing ones are not relevant to an individual.

<sup>5</sup>In fact, even more fundamentally, RT is a general theory of communication and cognition, but my concerns are restricted to the linguistic applications of RT.

<sup>6</sup>Those familiar with the framework might question whether the RT distinction between *implicature* versus *explicature* doesn’t introduce a third kind of meaning. However, these are crucially categories of *output* of the application of relevance-theoretic reasoning. They are the same kind of ‘content’ for the purposes of the present discussion since they are constituted by the same mechanisms.

and speech act content. And, crucially, they too reject the Gricean (and, more generally, moderate contextualist) association of linguistically encoded semantics with intuitively perceived truth-conditional meanings—i.e. with ‘what speakers say (assert, claim, state, etc.)’ by using given linguistic forms in given contexts. The rejection of this is so fundamental to C&L that they term it *the Mistaken Assumption*. They define this as follows (elsewhere, they call it the Speech Act Conception of Semantics; C&L to appear *a*, 2006*b*):

[The Mistaken Assumption:] A theory of semantic content is adequate just in case it accounts for all or most of the intuitions speakers have about speech act content, i.e., intuitions about what speakers say, assert, claim and state by uttering sentences. (Cappelen & Lepore 2005, 53)

What C&L apparently fail to recognise is that RT quite explicitly rejects this too (see further section 5). It may be that relevance theorists tend to express the idea in rather different terms—for example, ‘linguistically encoded semantics falls short of truth-conditionality’—but since truth conditions are conventionally identified with what the speaker is taken to be committed to (i.e. to have asserted, claimed or stated) by uttering some sentence, this is in fact just another way of expressing the very same point.

As mentioned above, C&L’s minimal semantic content is claimed to be propositional and this is a clear difference to RT’s claims about encoded meaning. As discussed in section 4, this does relate to certain genuine differences between the two approaches, but in other ways the importance of this contrast has been overstated. The fact that C&L’s particular notion of propositionality is significantly divorced from usual definitions causes a certain amount of confusion. More importantly, the propositionality or otherwise of minimal semantic content fails to distinguish it from RT’s ‘encoded meaning’ with regard to the role each plays in the sharing of content—and hence with regard to the evidence that C&L attempt to use against RT.

Given how closely aligned these two theoretical positions seem underlyingly to be, it is just as well for C&L that their own arguments against RT fail—as I show in the following sections.

### 3 C&L's criticisms of RT

#### 3.1 'Non-shared content' in reality and in theory

C&L's arguments rest on the idea that RT, like other radical contextualist approaches, is committed to the proposition that content is simply not shared (in some important sense). They derive this on the basis of quotations from relevance theorists like the following.

... communication can be successful without resulting in an exact duplication of thoughts in communicator and audience. We see communication as a matter of enlarging mutual cognitive environments, not of duplicating thoughts. (Sperber & Wilson 1986, 193)

Expanding on this, C&L claim that RT is crucially based upon a notion of *similarity* of content, instead of actual content-sharing (i.e. identity of content). In this they appear, on the face of it, to be justified by remarks such as the following:

An utterance, like any ostensive stimulus, usually licenses not a single interpretation, but any one of a number of interpretations with very similar import; provided the addressee recovers one of these, comprehension is successful, that is, it is good enough. (Carston 2004, 7)

There are at least two problems with C&L's characterisation of RT in this regard. First, it embodies claims that are stronger than anything claimed in RT (or in the above quotations): RT recognises that the addressee *may* not grasp an intended propositional meaning *p*; there is no claim that *p will* not be recovered. Individual relevance theorists might well vary over just how reliable they believe linguistic communication to be, but RT is at heart an attempt to explain just how commonly addressees do seem to recover intended propositional meanings (see the quotation from Carston below). It is rather important to the plausibility of RT that people should have reason to believe that they can recover each other's communicative intentions—otherwise it might not be worth trying—and if it were not the case that they regularly do so (or come negligibly close to it) they would probably soon find reasons *not* to believe any such thing. Hence RT in effect presupposes that, by almost any definition, content is regularly shared to a significant extent. However, the precise degree of reliability attributed to linguistic communication is not crucial to the present argument. This is because C&L's position rests on the idea that anything short of a guarantee of shared content is disastrous for communication,

so the validity of their arguments ultimately depends the truth or otherwise of this idea (which I come to in section 3.1.1).

The second, more important problem with C&L's position is that they imply that *similarity* is an active component of RT. It is not. C&L's criticisms look at content-sharing from the point of view of an omniscient third party. In this objective sense, intended meaning and recovered meaning may indeed be only similar. RT, however, is concerned with the addressee's understanding of what must be the intended meaning of an utterance (and how this understanding comes into being). It is this subjective assessment of intended meaning that drives RT as an explanatory framework and (as outlined below) this is also all that is required to understand the kinds of communication that C&L bring up in their arguments against RT.

It seems clear to me that Carston, in the quotation above, refers to similarity only from the point of view of an 'omniscient' analyst, without suggesting that this is part of the explanation of how communication is achieved. This is made quite plain by material from earlier in same paragraph:

The [presumption of an optimal rewards:effort ratio] provides a solution to the apparent problem [...] that the intended effect (the grasping of the communicator's meaning) cannot be independently observed to occur and then be recognized as desirable and presumably intentional. It provides a reliable, though by no means foolproof, means of inferring a speaker's meaning. (*ibid.*)

The actual work to which RT is put in communication, against which its success as a framework must be evaluated, has nothing to do with omniscient third parties, because such handy telepaths are not available for language-users to consult, and RT is a theory of what language-users do. Put another way, I do not read Carston here as claiming that the addressee says to himself, "I reckon that interpretation q is reasonably similar to the intended interpretation p and therefore communication has been successful". Rather, he says to himself, "Having employed the best means necessary [the Principle of Relevance], I calculate that the speaker must have intended to communicate q"—even if an omniscient third party might be able to identify that the speaker really intended to communicate the similar proposition p. It is only by failing to distinguish between these two ways of understanding the role of 'similarity' that one can imagine C&L's arguments to have any force against RT.

Some of C&L's misunderstanding seems to come from different notions of *successful communication*, a notion appealed to in the first quotation from Carston.

Here one might have some sympathy: there is a sense in which one would like to reserve this description for cases in which (at some crucial level) the intended meaning has been conveyed precisely. If we were to make this terminological decision, then we might be forced to conclude that, on the RT view, communication is rarely perfectly successful. However, even stating that communication may often be less than wholly successful in this sense is not to say that RT rules out its *ever* being wholly successful. Furthermore, it seems to me that a position that recognises the regular imperfection of communication reflects our everyday experience much more accurately than one that relies on infallible content-sharing to characterise human communication.

Just as importantly, as the above quotation from Sperber & Wilson states clearly, *successful communication* simply isn't defined in this way within RT. In RT, the goal of communication is to enlarge the interlocutors' mutual cognitive environment; that is (very roughly), communication is successful iff through it two people can tell, on the basis of the evidence available to them, that they have some more assumptions in common than they had before. This does not require the exact reproduction of thoughts (or of intended propositional meanings)—and, again, my intuition is that this is just as well: the exact reproduction of thoughts seems fantastically unlikely and even if it does happen we have no way of knowing it for sure, given the typical absence of telepathic semanticists from real communicative situations. Nevertheless, the process that just might get people to this happy situation does depend on the addressee's assumption that she can get at what the communicator must have meant to convey (so that it is worth trying so to do).

I do not mean to suggest by this that addressees are never aware that their understanding may be imperfect. We all know the feeling that we may or may not have fully grasped what the communicator meant (not least when reading the work of our fellow academics). This is in effect a failure of communication (or perhaps a problematically limited success, in Sperber & Wilson's sense of success) and it is only in such cases of partial failure that addressees consciously fall back on perceptions or hopes of mere similarity to the intended meaning. However common this may be, it is part of the *outcome* of relevance-theoretic reasoning in certain circumstances; it is not a part of that reasoning process. To suggest, then, as C&L do, that similarity constitutes a crucial part of the relevance-theoretic explanation of communication is to confuse RT, an explanation of what interlocutors do, with the possible *effects of applying* RT (as viewed by an omniscient third party).

This is important because any approach that actually took similarity to be an intrinsic part of the process of interpretation would quickly lapse into incoherence. C&L go to some lengths to show this (especially in C&L 2006*b*, to appear *a*), and

in so doing they merely succeed in very convincingly taking apart a straw man, at least where RT is concerned. Most of their discussion centres around the truth conditions of utterances that are reports of or reactions to another utterance. This gives grounds for suspicion from the outset, as it implies a semantic argument (i.e. one based in issues of abstract truth) made on the basis of essentially pragmatic issues (i.e. reports of or reactions to speech acts; see Gross 2006 for related argumentation). As counterarguments to RT, the arguments do not stand up any better in matters of detail.

Take the following argument from Cappelen & Lepore (to appear *a*). Imagine that Naomi has uttered a sentence *S* and someone else (call him Donald) reports that *Naomi said that p*, thus conveying that Donald believes Naomi to have conveyed the proposition *p* when she uttered *S*<sup>7</sup>. C&L then observe that, in a radically contextualist framework with no direct sharing of content, “the proposition attributed to Naomi will (typically) be similar to, but not identical to, the proposition that Naomi intended to communicate with her utterance” (to appear *a*, 5). As discussed above, a more appropriate wording would here replace *will (typically) be* with *may be*, but C&L’s subsequent remarks are more important: “A proponent of [such a framework] [...] has but two options: (a) she can say that reports like [*Naomi said that p*] are literally false, or (b) she can say that [*Naomi said that p*] can be true even though Naomi did not intend to communicate the proposition that *p*” (*ibid.*). Options (a) and (b) are what C&L proceed to call the Revisionist and Conservative strategies, respectively.

The so-called Conservative strategy (b) is what C&L (2006*b*) attribute to RT and other radical contextualist approaches, and they proceed there to show how incoherent such a strategy would be. It would, after all, amount to re-defining our very notions of truth and identity such that ‘similar’ propositions could be identified as being truth-conditionally equivalent. As C&L rightly point out, this would wreak havoc with all kinds of inferential procedures. They also point out (to appear *a*) that it would be essentially lacking in foundation, as no suitable definition of similarity could exist without a traditional notion of identity.

No matter; this strategy bears no relation to RT. I can only guess that C&L were led to think it is a conservative development of RT on the basis of relevance theorists’ references to *similarity* in successful communication, such as those in

---

<sup>7</sup>These annoyingly algebraic examples are useful because they allow us to avoid issues that are best sidestepped, such as the relationships between utterances, reports of utterances and propositional meanings, which are themselves a matter of apparently fundamental, and at any rate terminological, disagreement between C&L and relevance theorists. The above talk of ‘conveying a proposition’ is certainly shorthand for a series of processes from an RT perspective.

the quotations from Carston and Sperber & Wilson in section 3.1. But recall that the RT definition of successful communication makes no reference to objective evaluation of truth.

Returning to the example, Donald's reports of speech and attributions of belief are based on his subjective judgement of the identity of the proposition that Naomi intended to communicate. It is quite clear that the speech report is true as such *iff* Donald's judgement is right: i.e. *iff*  $p$  is (objectively) identical to Naomi's intended assertion. Otherwise, it will be, strictly speaking, false. But either way he can carry out the report sincerely and effectively. Ultimately, whether the objective truth or falsity of his statement really matters depends on a number of independent factors: just how far out Donald's judgement is and just how important it turns out to be in practical terms for Donald, Naomi and others to have understood one another very precisely. If it is important enough, it will doubtless become clear sooner or later just what was falsely assumed—their subsequent interactions and experiences will very likely provide them with ample evidence. If it is not so important, Donald and Naomi and any interested third party may well get by just fine, never knowing that he has made a report of her speech that is strictly speaking false in some way. A theory of the cognitive processes involved in communication has no responsibility to account for such extraneous factors or their effects (though RT does predict that the addressee will generally get at least close enough to the intended meaning for practical purposes, since he has a good set of principles to guide interpretation).

The same point applies to C&L's other examples in this vein, which involve phenomena such as attribution of belief based on previous utterances and explicit assessment of other people's assertions. RT does not, for example, claim that *Naomi believes that  $p$*  is equivalent to 'Naomi believes something similar to  $p$ ' or that *What you said is correct/wrong/outrageous* is equivalent to 'Something similar to what you said is correct/wrong/outrageous'. What it does claim is that because of the pervasive context-sensitivity of language, the speaker of *Naomi believes that  $p$*  may be, to a greater or lesser extent, wrong even if he believes he has heard Naomi express that  $p$ , and that sometimes a proposition judged to be correct, wrong or outrageous is not exactly the one that the original speaker intended to convey. Again, the degree to which this mismatch between speaker's thoughts and addressee's understanding actually matters is itself dependent on the nature of the context (as it develops). Considerations of objective truth (or of objective identity or similarity) simply do not enter into the process of communication and as such bear no relation to the only really important notion of content-sharing.

To summarise, C&L's 'similarity' argument is confused in several ways. It is an attempt to criticise RT as if it were a semantic theory, when RT is in fact a pragmatic theory (in the particular sense of a theory of the cognitive principles involved in communication). Furthermore (as discussed in section 4), it is self-contradictory for C&L to attempt to run a strictly semantic argument off data such as reports or assessments of utterances, since these all involve implicit reference to 'what has been said'. This clashes with C&L's rejection of the Mistaken Assumption, whereby the semantics of a natural language is not responsible for intuitions about 'what is said'. This kind of argument is therefore simply not available to C&L.

### 3.1.1 The R-strategy

The above rejection of the Conservative strategy clearly encapsulates a version of the so-called Revisionist one. Since (I maintain) this is the standard position within RT, there could hardly be a greater misnomer, so I shall refer to it as the R-strategy (for either *Relevance* or *Revisionist*, according to taste).

The R-strategy seems to me to be backed by strong intuitions and all kinds of anecdotal evidence. We can sincerely report the speech of others and, unknowingly, falsely convey that they meant to say something they didn't. Likewise, we can and surely do inadvertently attribute beliefs falsely as a result of misapprehending the intended significance of someone else's statements—even following very careful and extensive communication, as the histories of philosophy and linguistics surely demonstrate amply.

However, for C&L the obvious intuition is a different one: that people understand each other really rather precisely, rather a lot of the time, and that this is the basis of many kinds of complex behaviour<sup>8</sup>. This too is obviously true, in some sense. As C&L point out, people often follow instructions, laws and regulations that are conveyed linguistically and show appropriately co-ordinated responses across a range of individuals and contexts. People work together, often successfully, to achieve particular goals; this clearly requires shared understanding of those goals and effective discussion of potential solutions. Very generally, people act on what others communicate linguistically: we form beliefs and justify actions based on what others have said, and, relatedly, we hold others responsible for what they have said. Even relevance theorists behave as though language use

---

<sup>8</sup>At least, this is the line they push in their 2006*b*, to appear *a* articles. On the other hand, see C&L (2005, 206) for apparent recognition of the more problematic aspects of verbal communication.

is content-sharing: they write (more or less) carefully-worded tracts, attempting to specify precisely what they mean (see Cappelen & Lepore 2005, Chapter 9, for the charge that RT is self-contradictory in this way<sup>9</sup>). From such observations, C&L conclude that the R-strategy cannot be right: there must be some kind of content sharing that is reliable; language must indeed allow us to ‘reproduce thoughts’.

Putting together all the intuitive and anecdotal evidence of the above two paragraphs may seem to lead only to contradictions: people seem regularly to communicate with such striking effectiveness and with so much at stake that there is simply no room for error, yet we also feel that as communicators, as elsewhere in life, we are only human: we know that fundamental misunderstandings can occur. Though this may appear at first sight to represent a stalemate with regard to the arguments of the present article, it in fact constitutes an argument in favour of the R-strategy (and hence in favour of conventional RT).

To see why, consider the overall nature of C&L’s arguments. At first they make out that they stand in opposition to a complete lack of shared content in communication. This implies that they believe that successful communication is explicable as long as content sharing (in the sense of reproducing communicators’ thoughts) is *possible*. In this case, as I have argued, they should have no argument with RT: recall that Carston calls the principles of RT a “reliable but fallible” means to get at intended meanings (but not ‘hopelessly fallible’ or ‘inevitably inaccurate’). Therefore, if C&L see their observations as arguments against RT, they must either have misunderstood RT or believe that even the possibility of failing to grasp someone else’s intended meaning entirely cannot be admitted: i.e. content-sharing must be complete and guaranteed.

In fact, it seems that both of these things are the case. C&L really do seem to have missed the idea that the Principle of Relevance is meant principally to explain how addressees regularly do recover speakers’ meanings (even while predicting that they will sometimes fail to do this to a greater or lesser degree). Thus, when C&L refer to the basic RT process of inferentially developing the logical form of an utterance, they claim “We have no way to predict in advance which development of these logical forms various [addressees] will end up with . . . developments of these logical forms can be radically different and it would be a minor miracle if they were not” (to appear *a*, 18). To a relevance theorist, this is precisely a description of the problem that RT sets out to solve. One may argue about how well

---

<sup>9</sup>This argument is easily turned around: C&L’s often somewhat exasperated tone when discussing responses to their position suggests that they feel regularly misunderstood—i.e. they must have plenty of reasons to believe that content is not always shared very reliably.

RT does this, but one cannot criticise RT as though it fails to address this problem, when that is its very *raison d'être*. All the apparatus of RT, from the notions of manifestness and mutual cognitive environments to the cost-benefit calculations prompted by the Principles of Relevance, exists to provide an explanation of what would indeed be a minor miracle otherwise: the recognition of speakers' communicative intentions. Note that everyone, C&L included, requires an explanation of this, since it is uncontroversially required at least for the derivation of implicatures.

At the same time, C&L do seem to require complete and guaranteed content sharing—all their discussion points to this as being the only way to explain communication as we know it. I have already indicated that I find that the idea of guaranteed sharing of thoughts through language simply fails to reflect common experience. Furthermore, RT's recognition of a lack of guarantees merely puts communication into the context of very general facts about human processing of data about the outside world. We have only indirect access not just to other people's minds, but to the outside world as a whole: the evidence for all that we know or believe is filtered through our physical senses and through our abilities to process and reason over the data that they provide, both of which have been skewed in certain directions by our evolutionary history and are even then decidedly fallible. And yet we plainly have some pretty good ways of distinguishing more and less reliable judgements about how things are. The striking reliability of much linguistic communication is merely a part of this: we do not have direct access to our interlocutors' communicative intentions, but we plainly have some means of recognising and assessing evidence for what they are<sup>10</sup>.

Ultimately, the observations that C&L appeal to—linguistic communication as a basis for collective action, judging others' beliefs, etc.—logically show just the following: people *can* recognise each other's communicative intentions (i.e. share content) *well enough* to do all these things. RT not only accepts this, it suggests a mechanism by which it occurs.

Alongside the rather sweeping dismissal of RT (and with all radical contex-

---

<sup>10</sup>In this context, it is worth noting that RT has been used as a counter to the kinds of philosophy and, especially, literary theory that make sweeping claims about the 'impossibility of communication'. See, for example, MacKenzie (2002). In this connection, note that there is a certain parallel between literary 'Deconstructionists' and C&L. Both assume that linguistic communication is an 'all or nothing' business: either language allows meanings to be perfectly recreated or communication must be impossible. Deconstructionists argue that the latter conclusion must be reached, whereas C&L reason that since communication does seem to work pretty well on a regular basis, it must be the case that recovery of intended meaning is guaranteed.

tualism) discussed above, C&L (to appear *a*) do briefly attempt to explain why they think the particular mechanisms of RT fail to explain how communication is achieved. Their point (though wrapped up in much irrelevant discussion of ‘similarity’, as discussed above) is essentially that different contexts—both external circumstances and individual language users’ cognitive states—can be just too different. Attempting to adopt RT terminology, they state that “[what] you end with as a result of applying [relevance-theoretic reasoning], starting with the logical form of the sentence you’re interpreting [...] depends on what cognitive effects the utterance has on you, and that again on what beliefs (etc.) that the interpreter has” (*sic*).

This is a misrepresentation of RT, implying that the assumptions of different interlocutors are quite divorced from each other and subject to unconstrained, idiosyncratic variation; that there is no basis for relating one’s own assumptions to those of another language user. This ignores the central role of mutual manifestness in RT. Far from depending on just any ‘beliefs, etc.’ that an individual addressee may entertain, interpretation is based on that which interlocutors have good reason to believe that they share (something that we need for many purposes other than linguistic communication). This always exists, though it will be richer in some cases than in others—a fact that affects the way one expresses oneself in different contexts (even in C&L’s example of their writing for an unidentified reader they evidently have a *kind* of reader in mind; one with certain assumptions that they can rely on). C&L’s description of the ‘cognitive effects an utterance has on you’ suggests a kind of free association, the quirky sparking of an individual’s imagination. But RT is concerned with the ways we recognise just those effects that are intended by the speaker of an utterance. Once again, C&L’s attempted criticism of RT works better as a description of the kind of problem that RT sets out to solve—in other words, how things are if one does *not* assume something like RT.

At this point, a supporter of C&L might retreat to another position, claiming that even if RT hasn’t been refuted, either approach is still possible. It might be conceded that RT offers one kind of explanation for the recognition of others’ communicative intentions (in contradiction to C&L’s 2005, Chapter 8 claim that radical contextualism “makes communication impossible”), but also noted that C&L’s notion of content-sharing directly through language offers another kind of explanation, and arguably a much simpler one. This might be the case, were it not for the fact that C&L’s notion of content-sharing is self-contradictory, as shown in the following section.

## 4 Relevant kinds of content

So far I have outlined several problems with C&L's arguments regarding shared content which in effect involve misreadings of RT. Their arguments also suffer from a fundamental problem of internal incoherence. This concerns the relevant understanding of *content*. My argument here runs as follows:

- C&L argue that communication, as we know it from various examples of human behaviour, depends upon guaranteed content sharing.
- The only kind of content that could be infallibly shared, on any view—and the only kind that has one of the defining properties of C&L's 'shared content'—is that which RT calls linguistically encoded meaning. C&L's minimal semantic content is similar to this in important ways and it is indeed their proposal that this is the content that is shared.
- However, the kind of content that is shared in C&L's examples of human behaviour is not this, but rather what they call Speech Act Content, which no-one claims to be determinate.
- So, by C&L's own arguments, communication should be impossible—in just the ways that they observe it not to be.

We have already seen how C&L urge the first point in this list (the first premise in my argument). The second and third points are elucidated in the following sections.

### 4.1 Minimal semantic content as shared content

C&L place three key requirements on their level of minimal semantic content:

- i. that it be the meaning that is constant across all uses of a given linguistic form, in all contexts
- ii. that it be propositional
- iii. that it be the kind of shared content on which are based such activities as the evaluation of assertions, reporting of indirect speech acts, justification of subsequent actions, and so on

In the following section, I show how (iii) cannot be correct. In section 4.2, I discuss (ii). This is the source of a substantial difference between C&L’s minimal semantic content and RT’s encoded meaning. However, as argued below (and at length by Carston 2006, Korta & Perry to appear, Borg to appear), it is ultimately incompatible with (i)—at least given C&L’s notion of propositionality. Furthermore, C&L’s heuristics for minimal semantic content fail to identify necessarily propositional kinds of content. In this section, I concentrate on the implications of defining minimal semantic content in terms of (i).

That minimal semantic content should be constant across all contexts is stated explicitly by C&L (2005, 143) and is implicit in both their arguments about the nature of shared content and in the way they state minimal semantic content truth-conditionally. When C&L speak of truth conditions, they do not refer to the common (moderate contextualist) conception of ‘what makes the propositional intuitively expressed by a given utterance true’, but rather something that is indeed considerably more minimal. (1) is an example<sup>11</sup>:

- (1) ‘Steel isn’t strong enough’ expresses the proposition *that steel isn’t strong enough* and is true iff steel isn’t strong enough” (C&L 2005, 61)

Though this kind of statement appears less than dazzlingly enlightening, C&L insist that this is the level at which truth conditions should be stated, arguing that it is the business of metaphysics, not semantics, to determine conditions on truth any more precisely<sup>12</sup>. The point for current purposes is that (1) shows how C&L’s strictly semantic argumentation—in common with their essentially communication-based ‘shared content’ arguments—is committed to the idea that the semantic content of a given sentence is just whatever remains common to all its meaning across all contexts of utterance. There is no other way to interpret statements like (1).

By this definition, C&L’s minimal semantic content is entirely equivalent to the notion of encoded meaning in RT. This is necessarily the case, since encoded

---

<sup>11</sup>As stated above, I will ignore for now the fact that C&L assume there is a *proposition* ‘that steel isn’t strong enough’. This brings problems once unequivocally indexical expressions, such as pronouns, are involved; see section 4.2.

<sup>12</sup>*Contra* the avowed radical contextualist Travis (2006), I believe that there is real substance to this point and that it is essentially compatible with radical contextualist thought. RT work on *ad hoc concepts* (see Carston 2002, Chapter 5) could be seen as arguing something roughly equivalent from a linguist’s point of view. In arguing, in effect, that the likes of (1) just aren’t proper truth conditions, Travis (and arguably also Bezuidenhout 2006) stray close to endorsing a moderate contextualist stance.

meaning is by definition just that which is conveyed by a given linguistic form, irrespective of context. Note that to say that some meaning is encoded is entirely neutral with respect to questions such as whether it has other properties, such as being propositional. Thus, RT's notion of encoded meaning is equivalent to minimal semantic content as long as the latter conforms to (i) (as it must, to carry any of C&L's arguments), but is free of the problems that (ii) brings, as discussed below.

Given this, it is unsurprising that C&L's discussion of the psychological and communicative roles played by minimal semantic content fails entirely to distinguish their position from that of RT. In an entire chapter dedicated to these issues (2005, Chapter 12) they merely reach the irrelevant conclusion that RT's arguments against Gricean theorising would not constitute an argument against their notion of 'semantic minimalism' (C&L 2005, 180; this is, of course, consistent with my argument that the similarities between C&L's position and RT are substantial and outweigh the differences). In the same chapter, C&L show that they are aware that parallels between their position and RT might be drawn and are (for some reason) very keen to deny their validity. They draw up a list of five differences between the two approaches (2005, 181). All but one of these rest upon the supposedly crucial role of the propositionality of minimal semantic content<sup>13</sup>. However, this plays no role in C&L's discussions of communication and brings significant problems at other levels, as outlined below.

## 4.2 Minimal semantic content as propositional content

Differences between 'minimal semantic content' and 'encoded meaning' arise because of C&L's second requirement (or assumption) regarding minimal semantic content: that is propositional. Tellingly, this is also where both logical and empirical problems enter into C&L's story.

C&L insist that the propositionality of minimal semantic content is crucial to its psychological and communicative role (e.g. 2005, 181, and the reply to Korta & Perry to appear in C&L to appear *b*), but it is quite unclear why this should be, given that (as outlined above and shown below) all the communicative phenomena that C&L invoke depend on the *existence* of a crucial level of shared content, not on its logical properties (Borg, to appear, suggests some reasons why philoso-

---

<sup>13</sup>The other one concerns C&L's notion of *Original Utterance Centrism*, which I cannot discuss here for reasons of space. This is not immediately relevant to current concerns, except that it once again demonstrates the inappropriate adoption of the semanticist's omniscient third party perspective in the discussion of *what is communicated* on a given occasion.

phers might care about the propositionality of semantic content, but notes that these are independent of the analysis of communication). It is of course necessary that some linguistically conveyed meaning be propositional, as shown by activities like assessing the truth or plausibility of what someone else says. However, I show below that this cannot be the same kind of meaning as that which is constant across different contexts (and is thus ‘shared’ in C&L’s sense) and as such this observation provides no support for C&L’s assumption that minimal semantic content must be propositional.

C&L’s self-imposed commitment to the propositionality of minimal semantic content forces them to make this content less minimal than it could be, as Borg (to appear) and Korta & Perry (to appear) point out. It also introduces a difference between their conception of minimal semantic content and RT’s encoded meaning, as Carston (2006) stresses. In order to stick to what they recognise to be propositional meanings, C&L have to assume that minimal semantic content contains no unsaturated indexicals or ambiguities. The problem is that this immediately makes minimal semantic content something other than that which is shared across different contexts: the saturation of indexicals is a context-dependent process. C&L note that this means their minimal semantic content “cannot be characterized *completely* independently of context” (2005, 143; emphasis in original), but do not explain how this is to be squared with the proposed role of minimal semantic content as infallibly shared content.

Here C&L seem to be following the convention in semantics of simply assuming that there exists some self-contained process of assignment of reference to indexicals and that the semanticist has no responsibility to characterise what this is or how it relates to other context-dependent areas of linguistic interpretation (see Larson & Segal 1995 for a textbook non-exposition of this idea, for example). This is curious, given that C&L’s general concern is meant to be precisely the matter of context-dependence in language. It is particularly strange that C&L should include the referents of overt indexicals in that meaning that is infallibly shared, since this is one of the more obvious sources of misinterpretation in language use (given an utterance of *He’s ready*, surely a response like *Who? Mark or Paul?* is at least as likely as *Ready for what?*, yet according to C&L’s position, the answer to the former question to be ‘shared content’, while the answer to the latter is not).

In this sense, then, C&L’s particular formulation of minimal semantic content is insufficiently minimal to achieve one of the key roles that they assign to it. I refer the reader to the above-cited works for more fully developed and philosophically grounded versions of this argument; my current concern is with the degree

to which the propositionality of minimal semantic content plays any genuine role in C&L's arguments regarding content-sharing—and the degree to which propositionality is in any case demonstrated in their heuristics for minimal semantic content. If propositionality is not necessary in either of these ways, the obvious step for C&L or those who otherwise have sympathy for their arguments is to adopt the RT notion of encoded meaning, which provides the logically necessary level of some infallibly shared content without the problems associated with C&L's propositional version.

In fact, no part of C&L's argumentation on the role of minimal semantic content in communication actually seems to depend on its propositional status. Consider the following quotation, which is typical of C&L's reasoning in this area:

How is it that we can understand what was said by an utterance of [*Peter's duck is brown*] when that utterance took place in a context radically different from ours (and we know little or nothing about that context)? The answer should be obvious by now: We can always understand part of what the speaker said, namely *that Peter's duck is brown*. (C&L 2005, 205)

The argument here turns upon there being some constant meaning attached to a given linguistic form (such as a sentence). It does not depend upon this constant meaning having any particular properties, such as being propositional<sup>14</sup>. Consequently, the role of minimal semantic content, as envisaged by C&L themselves, does not distinguish it from the encoded meaning that RT makes use of, at least as far as content-sharing is concerned.

To summarise this section, C&L's minimal semantic content is torn in two by conflicting requirements that they place on it. The first is that it should be context-independent meaning. Any level of meaning that genuinely fulfils this criterion is of necessity equivalent to RT's encoded meaning, since this is definitive of encoded meaning. As such, C&L have no room to criticise RT for lacking minimal semantic content, on this definition. The second is that minimal semantic content must be propositional (in C&L's minimal sense). While playing no discernible role in the explanation of communication, this conflicts with both the first requirement and, as shown in section 4.3, the ways in which C&L attempt to identify such content.

---

<sup>14</sup>In case it seems that being the complement clause of *say* guarantees propositionality in any pertinent sense, see section 4.3.

### 4.3 Tests for minimal semantic content

C&L (2005, Chapter 7) offer certain tests for minimal semantic content, which are important in numerous ways. They represent C&L's concrete methodological alternative to working only with the moderate contextualist's intuited propositions and they supposedly constitute C&L's most substantial objection to radical contextualism<sup>15</sup>. Moreover, they provide the only concrete evidence for propositionality in minimal semantic content. These tests are close relations of some of the tests for shared content mentioned above: again, the basic strategy is to examine reported speech, thus turning a judgement about the *content* of an original utterance into a judgement about the *truth* of a second utterance, the speech report.

C&L's primary concern in presenting these tests is to take on contextualism as a whole. They are keen to point out that the use of speech reports distinguishes uncontroversial, overt indexicals, such as pronouns, from other cases of alleged context-sensitivity. Thus, (2b) would be false as a report of (2a) because of the context-sensitivity of the pronoun *I* (Cappelen & Lepore 2005, 89).

- (2) a. Utterance made by George Bush, June 3, 2003: 'I wasn't ready yesterday.'
- b. Utterance made by Ernest Lepore, June 5, 2005: 'Bush said that I wasn't ready yesterday.'

In contrast, had Bush uttered a sentence like *John is ready*, it could be truthfully reported by Lepore with the utterance *Bush said that John is ready*. This is despite the fact that *John is ready* intuitively falls short of specifying full truth conditions (it invites the question *Ready for what?*), so is context-sensitive by contextualist reasoning.

C&L's tests thus undoubtedly demonstrate that items like pronouns (what C&L call the Basic Set of indexicals) are different to other linguistic material. This is problematic for the kind of moderate contextualist who would deal with the 'incomplete truth conditions' of strings like *John is ready* by positing covert indexical elements, but not for genuinely radical contextualists such as relevance theorists. By definition, radical contextualist approaches take the influence of context to be pervasive. Hence it is perfectly consistent with a radical contextualist position to accept that there are some linguistic items that specially encode a particular relationship to context (true indexicals), while others are prey to the influence of

---

<sup>15</sup>C&L's other objections to radical contextualism (2005, Chapters 8 and 9) amount to a version of the case against RT already discussed above.

context in other ways<sup>16</sup>. C&L’s argument regarding true indexicals and context-sensitivity therefore only has force as an objection to moderate contextualism, even though they present it as an objection to the radical kind<sup>17</sup>.

Still, perhaps these tests do demonstrate something important: the key difference between C&L and RT, which is C&L’s claim that encoded meaning is propositional. How else, the argument goes, could we make judgements of truth and falsity about indirect speech reports? Intuitively, when we say *Nina said that John is ready* (unlike *Nina said, “John is ready”*) we express a proposition concerning Nina’s having committed herself to something—something that is true only under certain conditions. Therefore, it seems that any utterance which can be indirectly reported in this way itself expresses a proposition, not just a ‘propositional radical’ or any other form of sub-propositional encoded meaning.

To give another perspective on the same idea, C&L suggest imagining two contexts C1 and C2 in which Nina has pronounced the words *John is ready*. Even if these contexts are sufficiently different that John was ready for quite different things, we can say that the sentence *In both C1 and C2, Nina said that John is ready* is true. C&L take this to be an argument against radical contextualism: if *ready* had to be enriched in order to create a proposition, then in this sentence the complement of *said* would have to express two distinct propositions simultaneously. This is incompatible with the existence of a single truth value for the whole sentence.

C&L’s reasoning here rests on a crucial empirical question that they do not investigate: do indirect speech reports consistently work this way? In particular, can we really make reliable judgements of truth and falsity about such reports? Note that this consistency is crucial to C&L’s arguments; if the answer to the question is no, we can no longer assume that the complement of *said* in such cases is necessarily propositional. In that case, being the complement of *said* is simply not a test for propositionality—meaning that indirect speech reports cannot be used to argue for the propositionality of encoded meaning that putatively distinguishes semantic minimalism from RT<sup>18</sup>.

---

<sup>16</sup>See Atlas (to appear) for a careful elaboration of the general point that recognising a case of context-sensitivity does not entail a claim of indexicality of any kind.

<sup>17</sup>As illustrated in section 5.1, it may be that professed radical contextualists sometimes behave in practice more like moderate contextualists (in positing covert indexicals and the like just to account for context-dependent changes in propositional meaning), but the present argument is about radical contextualism as a philosophical idea, not about individuals’ practices.

<sup>18</sup>Therefore in this case C&L therefore cannot argue that they merely have to show that the appropriate judgements *can* be made in at least one case, as they argue in Cappelen & Lepore

There are various ways in which indirect speech reports of the kind that C&L rely on can fail to report a single proposition. For example, being the complement of *said* does not force disambiguation of homonyms. I can truthfully utter the sentence *In C1 and C2, Nina said that John went to the bank* even if I know that Nina was referring in C1 to the financial *bank* and in C2 to the watercourse *bank*. Of course, I might well continue . . . *but in C1 she meant that he went to the HSBC whereas in C2 she meant that he went to the edge of the river*. Significantly, one may easily do this with regard to C&L's example also: *In C1 and C2, Nina said that John is ready—but in C1 she meant he is prepared for his exam, while in C2 she meant he has his raincoat on*. I can see no difference in the form of these utterances that would allow us to treat them as identifiably different grammatical phenomena (the one an ambiguity, hence potentially expressing different propositions under any analysis, and the other a single proposition in C&L's sense). Therefore, unless C&L want to argue that an un-disambiguated homonym contributes to the creation of a proposition, they cannot maintain that indirect reports necessarily contain propositions in their complement clauses.

Similarly, C&L's test fails to disambiguate use/mention ambiguities: consider (3):

- (3) (It is true that) in C1 and C2, Elspeth said she didn't want to work in Human Resources  
 . . . but in C1 she meant she wanted to work for a department called Personnel, while in C2 she meant that she didn't want to work in that area at all.

I doubt that C&L would wish to claim that any proposition exists before such an ambiguity is resolved, but they would have to do this in order to maintain the supposed significance of their indirect speech report test. Again, note that C&L do not have the option here of simply dismissing such cases as being intuitively different to their examples. Their avowed purpose in introducing linguistic tests is to get away from reliance on intuitions (2005, 87); as such, the tests they propose in their place cannot themselves depend on intuition but must be purely linguistically defined and must have consistent effects in order to convince.

It seems that indirect speech reports are in a sense not indirect enough to do the job that C&L ask of them. They cannot be relied upon to identify a propositional minimal semantic content and as such they do not demonstrate that minimal semantic content is necessarily propositional. Without this demonstration, we are essentially left in RT territory: we know there is encoded meaning, but we can't

---

(2006a).

prejudge the nature of it for any given linguistic expression (and, consequently, inferential pragmatic processes may in principle influence all levels of meaning beyond this).

So much for the theory of minimal semantic content. Where C&L's (2006*b*, to appear *a*) empirical arguments are concerned, there is an equally important problem: the role of minimal semantic content simply does not extend to the functions of shared content that are set out in these arguments. These must involve a different kind of content altogether.

#### 4.4 The content that is shared

Consider again the kinds of observations that C&L claim make shared content indispensable: co-ordinated action, collective deliberation, linguistic communication justifying beliefs, holding people responsible for what they say, and so on. Is minimal semantic content capable of supporting such things? If so, then an engineer who says *Steel isn't strong enough* is held responsible just for that proposition (if we call it that) which is expressed by any and every utterance of this sentence. If the roof collapses, it should be considered reasonable of the engineer to point out that *I just said steel is strong enough; I never said strong enough to support the roof*. Such utterances would usually be described using words other than *reasonable*. The example, crass as it may seem, is not a trivial one in the context of C&L's discussion of content sharing: it shows that if 'being held responsible for what one says' is evidence for content sharing (as they claim it is), then the relevant kind of content is not C&L's minimal semantic content.

Similar arguments apply to C&L's other examples of content sharing: if Beth says to me *The apple is red*, I will not be tempted to conclude that 'Beth believes [the minimal proposition that is true iff the apple is red]'—something that, as C&L recognise, could cover a huge number of possible conditions, in different contexts (e.g. the apple is wholly red, partially red, red on the inside, painted with red stripes, mostly a colour that would otherwise be called brown but therefore 'red' as opposed to being a green apple, etc.). I might well, however, conclude that 'Beth believes that the apple is red in the particular sense that it is partially red on the outside' (for example). Similarly, if I am part of "a CIA task force concerned with whether Igor knows that Jane is a spy" (C&L to appear *a*, 9), I will assume that we are 'collectively deliberating' about some rather more specific proposition than the minimal semantic content of *Igor knows that Jane is a spy* (it probably won't concern us if Igor has figured out that Jane is revealing the secrets of one knitting circle to the members of another, as long as he hasn't

noticed that she's also a CIA spook—but both of these situations are compatible with this minimal semantic content). As a final example, consider again C&L's evidence from people's assessments of others' assertions. The Speech Act nature of the relevant 'shared' content here is shown by the fact that such assessments can be made about implicatures just as easily as about what is taken to be directly asserted, as shown in (4):

- (4)       A: Is John really such a bad cook?  
          B: Well, he's English.  
          A: That's outrageous / a fair point.

It is clear that minimal semantic content is not what must be shared in order to facilitate these activities. What is needed is, minimally, some notion of 'what is said' by a given utterance. As we have seen, C&L's rejection of the Mistaken Assumption makes quite clear that this is not a significant part of their theory. Instead, 'what is said' is merely a part of Speech Act (henceforth SA) content. Therefore, what is necessarily shared here is SA content.

Once this is established, all of C&L's arguments regarding RT and content-sharing simply fall apart. If the basis of their examples of the necessity of content-sharing is the sharing of SA content, then these examples fail to distinguish in any way between the position of RT and that of C&L themselves. C&L are quite explicit on the fact that, for them as in RT, the recovery of SA content is indirect, context-dependent and indeterminate<sup>19</sup>

On our view, any utterance succeeds in expressing an indefinite number of propositions. One of these, the proposition semantically expressed [=minimal semantic content], is easy to grasp. Others are extremely hard to access and there is no reason to think that any *one* person can ever grasp all that was said by an utterance  
(C&L 2005, 206; emphasis in original)

Also just as in RT, C&L argue that SA content is in part determined by the same context-dependent mechanisms that generate implicatures (this is their eighth principle of Speech Act Pluralism; 2005, 204). Since C&L's SA content encompasses

---

<sup>19</sup>RT would claim not that encoded meaning is one of a number of propositions but that an indefinite number of propositions may be constructed on the basis of it, as it interacts with context. Again, the difference between the two approaches effectively reduces to the issue of putative propositionality at the minimal level of meaning, which we have seen to be neither well-supported empirically nor necessary to C&L's own mode of analysis.

what most (moderate contextualist) semanticists would call the propositional content of a sentence, this is a striking convergence with RT on a position that opposes dominant assumptions.

To summarise, the basis of collective deliberation, inferences of belief, and so on is demonstrably SA content, not minimal semantic (or encoded) content. C&L and RT agree that SA content is not infallibly shared. It follows that the degree to which these activities involve ‘sharing’ of content is counter-evidence to RT only to the same extent that they are counter-evidence to C&L’s own position. Fortunately for both, we have seen that the arguments made around these phenomena are in any case misdirected: perfect content-sharing in this sense is not essential to communication or to the theory of communication.

Given their position on the nature of SA content, C&L in effect presuppose an approach that is at least broadly like RT. Ostensibly, they avoid committing themselves to any particular account of how SA content is derived, claiming that “there can be no *systematic theory* of speech act content” (2005, 190; emphasis in original) and thereby washing their hands of the responsibility to formulate one. Again, I see a failure to distinguish objective, external characterisations of individual acts of communication from characterisations of the internal competences that allow people to make the necessary subjective judgements to communicate effectively. While a fully formalisable mechanism showing the derivation of SA content is most likely quite impossible, one can theorise about the principles that people employ when required to derive it<sup>20</sup>.

## 5 RT and the radical/moderate divide

C&L’s distinction between moderate and radical contextualism is an important one. On the other hand, as should now be clear, the reason their arguments against RT fail is essentially because they fail to recognise significant similarities between true radical contextualism and true semantic minimalism. These are quite logical and predictable: the more minimal one’s notion of semantic content is, the more work there is to be done by non-semantic, context-dependent processes. C&L’s hostility to radical contextualism is therefore rather surprising. I believe that it is based upon certain points of confusion as much as on substantive disagreements

---

<sup>20</sup>C&L (2006*b*) do note that their position requires answers to questions such as “How is speech act content determined?” and “how interpreters focus sometimes on one aspect of the speech act content, sometimes on another”. It is unclear why they fail to see RT—and indeed the Gricean and post-Gricean pragmatic traditions as a whole—as an attempt to answer exactly such questions.

of principle. One of the key points of confusion is C&L's idea that both moderate and radical contextualists (including relevance theorists) subscribe to the Mistaken Assumption. By their own definitions, no-one who did would actually be a radical contextualist.

The only reason C&L give for lumping together moderate and radical contextualists in this respect (apart from a single quotation from the supposed radical contextualist—but not relevance theorist—Charles Travis) is their common use of what C&L call Context Shifting Arguments (CSAs). This refers to the strategy of showing that what is intuitively asserted by some sentence varies from context to context and on the basis of this claiming that the semantics of the sentence in question is context-dependent. C&L believe that a commitment to the Mistaken Assumption follows from the use of CSAs:

It's the only way to make sense of their extensive (mis)use of CSAs. If you did not believe in some version of [the Mistaken Assumption], why would you care about the intuitions that speakers have that an utterance of 'There are no French girls' can be used to say (assert, claim, state) that there are no French girls in Room 401? Why would a semanticist hold that piece of information relevant? Why, more generally, would a semanticist think any of the intuitions appealed to in CSAs are relevant? (C&L 2005, 54)

There is a straightforward answer to C&L's rhetorical question, from the radical contextualist's perspective: the intuitions appealed to in CSAs are relevant to the debates in which we are engaged just because almost everyone else assumes them to be theoretically important. Radical contextualists' use of CSAs can be viewed as implicitly doing the the very same job as C&L's own (2005, Chapter 5) arguments about moderate contextualism: they make a conditional argument. *If* you believe that 'what is said (asserted, claimed, stated, etc.)' is the appropriate level at which to state truth conditions, then you will soon find that you cannot maintain a direct relationship between encoded meaning and truth-conditional meaning—contextual influence is simply too pervasive. In other words, the Mistaken Assumption is a premise used in a *reductio* of moderate contextualism by both C&L and radical contextualists.

Another way to understand the issue is to recognise that what C&L call *the* Mistaken Assumption in fact has two distinct parts. The first is the assumption that 'what is said (asserted, claimed, stated, etc.)' is the appropriate level of meaning at which to invoke truth conditions. C&L deduce that radical contextualists hold the Mistaken Assumption on the grounds that they seem to subscribe to this first

assumption (this is clearly the nature of their commentary on a quotation from Travis, for example; C&L 2005, 56). However, this amounts to a mere terminological point unless one accept the second part of the Mistaken Assumption: that the semantics of linguistic expressions should account for this level of meaning directly. It is in the very nature of radical contextualism to deny this second assumption and, as such, to deny the Mistaken Assumption. It is somewhat immaterial whether or not one thinks that ‘what is said’ might be spoken of as ‘the truth conditions of an utterance’ if one’s express purpose is to show that this level of meaning is not directly related to the semantics of the linguistic expressions uttered.

If C&L’s distinction between moderate and radical contextualism has any substance, then what sets apart the radical kind is surely the claim of the *generality* of contextual influence—or rather (to get to the heart of the RT perspective) the pervasiveness of the influence of extra-linguistic processes on the interpretation of linguistic forms. This means that truly radical-contextualist arguments are intrinsically incompatible with the crucial part of the Mistaken Assumption, which depends upon contextual influence on ‘what is said’ being restricted in such a way that a direct connection between it and encoded meaning can be maintained. It follows that any argument that involves the Mistaken Assumption is by definition a moderate contextualist argument.

Other misconceptions may have arisen, ironically enough, from the very fact that relevance theorists and C&L have a common enemy in moderate contextualism. The problem is that they approach this common enemy from opposite perspectives (in a way that is doubtless related to the origins of RT in linguistics and anthropology, while C&L are philosophers) and as such may fail to realise just how much of a common cause they have<sup>21</sup>. Thus, relevance theorists tend to push the notion that contextual influence is pervasive, in contrast to the limited, grammatically triggered contextual influence allowed for in the dominant moderate contextualist paradigm. As part of this, considerable emphasis is placed on the idea that propositional meanings are partly derived by context-dependent inference. Since at this stage the point is to take on moderate contextualism, the propositional meanings referred to here are those intuited truth-conditional forms that moderate contextualists deal in.

---

<sup>21</sup>The above-cited eighth principle of Speech Act Pluralism (2005, 204), one of the few places in which C&L explicitly confront moderate contextualist assumptions from the point of view of how much contextual influence on meaning *is* necessary, reads like an excerpt from an exegesis of RT ideas—or indeed of those of several other past and present contextualist analysts, as Atlas (to appear, 14) notes.

C&L's notion of proposition is quite removed from that of moderate contextualists (they put a good deal of effort into arguing that what moderate contextualists consider to be sub-propositional is for them fully propositional; C&L 2005, Chapter 5). Thus, we clearly have two separate notions—we could call them Propositionality<sub>CL</sub> and Propositionality<sub>MC</sub>. When a relevance theorist denies that encoded meaning is propositional (in line with C&L's definition of radical contextualism), this refers to Propositionality<sub>MC</sub>.

As we have seen (in section 4.2), Propositionality<sub>CL</sub> is not ultimately compatible with RT either. However, this is only because of C&L's line on the saturation of recognised indexicals. This is a genuine and important difference, but nothing like as important as that between moderate contextualism and either C&L or RT. As long as the discussion concerns examples without recognised indexicals, such as (1), there is little or no conflict between RT and C&L, because both recognise a level of meaning that consists of just whatever is common to all uses of a given linguistic form. If this is Propositionality<sub>CL</sub>, then (where these examples are concerned) RT logically *cannot* be arguing for contextual influence in the creation of Propositionality<sub>CL</sub>. Meanwhile, C&L recognise that there is contextual influence in the creation of Propositionality<sub>CL</sub> if the sentence does contain a recognised indexical (2005, 143). Thus C&L are in fact in total agreement with RT on the issue of which levels of meaning are subject to contextual influence: everything above and beyond encoded meaning, including the saturation of indexicals. Genuine differences between the two begin with C&L's idea that there is a single, minimal, theoretically and psychologically significant level of meaning that mixes context-independent and context-dependent aspects of meaning. However, this point of divergence is not the logical or ostensible basis of C&L's arguments against RT, which instead rest upon the false notion that RT sees contextual influence where C&L see none.

## 5.1 Moderation as a cardinal sin

Having thrown a lot of criticism at C&L for their take on RT, it is time to give due recognition to their analysis of pivotal assumptions in semantic/pragmatic theory. The nature of my argument has been that RT is immune to the more coherent parts of C&L's criticism just because RT is truly radical (i.e. thorough) in its contextualism. It follows that any approach that strays into moderate contextualism becomes a legitimate target for their minimalist objections. As such, I believe that C&L's notions of the radical/moderate divide and of the Mistaken Assumption can serve as important points of reference for the assessment of both theory and practice in

linguistic analysis (whence my ready adoption of these terms in the present work).

In many ways, the dominant moderate contextualist (e.g. Gricean) view sets the agenda in linguistic semantics and pragmatics. This can skew discussions, leading towards excessive concerns over ‘truth-conditionality’ (i.e. Propositionality<sub>MC</sub>), even sometimes among those analysts who clearly fall on the ‘radical’ side of C&L’s distinction<sup>22</sup>.

Relevance theorists are not immune to the tug of moderate contextualism. One case in point is Carston’s (1998) analysis of scalar inferences involving cardinal numerals (surprisingly, since Carston’s work is typically characterised by what is in effect thorough argumentation for genuine radical contextualism). Carston’s main aim here is to argue against the (unfortunately still widely assumed) neo-Gricean analysis of numerals (originally due to Horn 1972), whereby they are taken to encode ‘at least *n*’ readings and take on ‘exactly *n*’ readings only by implicature. She produces many convincing arguments to this effect, but moves on to propose a categorical distinction between numerals and other expressions that trigger what would traditionally be termed ‘quantity’ inference (after Grice’s 1975 Maxims of Quantity)—including upper-bounded readings of non-numeral quantifiers like *many* as well as inferences of ‘exhaustive listing’ of clearly ‘non-scalar’ objects.

The problem with this is the nature of the reasoning that Carston employs. On the basis of examples like (5) and (6) (originally due to Horn 1992), she argues that the upper-bounded (‘exactly two’) reading of numerals constitutes part of the truth-conditional meaning that the speaker is committed to, while the upper-bounded reading of non-numerals (e.g. ‘many but not all’) is mere cancellable implicature<sup>23</sup>.

- (5)       A: Do you have two children?  
          B<sub>1</sub>: No, three.  
          B<sub>2</sub>: ?Yes, (in fact) three.
- (6)       A: Are many of your friends linguists?  
          B<sub>1</sub>: ?No, all of them.  
          B<sub>2</sub>: Yes, (in fact) all of them.

---

<sup>22</sup>It is open to question whether all of those labelled radical contextualist by C&L consistently warrant the label, according to the criterion of fully rejecting the Mistaken Assumption. I continue to limit the discussion here to RT, which at least in principle does fulfil this criterion.

<sup>23</sup>Geurts (1998) rightly calls into question the use of *yes/no* responses as markers of truth-conditionality (by any definition), but my focus here is the broader structure of Carston’s argument.

Carston then proposes a ‘sense-general’ lexical semantics for numerals—i.e. one which is systematically underspecified for ‘at least’, ‘exactly’ and ‘at most’ readings (see also Atlas 2005). She proposes a rough formalisation of this as [X[THREE]]. As Carston explains,

This representation overtly requires that material be supplied pragmatically to instantiate the variable X; that is, the necessity of a process of pragmatic enrichment is signalled in the logical form (semantic representation) of the utterance. (Carston 1998)

This introduction of a ‘contextual variable’ is recognisably the formal strategy of many moderate contextualists<sup>24</sup>. Indeed, Carston explicitly draws on moderate contextualist proposals for quantifier domain restriction<sup>25</sup>. This might be viewed as a mere formal detail; however, it directly reflects the reasoning that underlies Carston’s analysis. This reasoning is definitively moderate contextualist: Carston proposes a certain kind of encoded semantics for the express purpose of accounting for context-dependent effects in intuitions of truth-conditional meaning. This is quite simply the Mistaken Assumption in action.

There clearly are significant differences between numerals and other expressions. Furthermore, contrasts like that between (5) and (6) are in principle interesting and worthy of explanation. But in RT the issue should never be framed in terms of a categorical distinction between items that affect truth conditions and those that don’t<sup>26</sup>. There are many inherent and fairly uncontroversial properties of numerals, at both the semantic level and in likely conditions of use, which can be argued to account for contrasts like (5) versus (6)<sup>27</sup>. Furthermore, such con-

---

<sup>24</sup>This is not to say that ‘hidden elements’ in logical form can never be justified under RT assumptions. Were there some other form of evidence—e.g. syntactic—for such an element, this would be a different matter, to be judged on a case-by-case basis. What is not compatible with RT is specifically reasoning from ‘effects on truth conditions’ to the existence of hidden elements.

<sup>25</sup>These are subsequently defended at length by the ‘hidden indexicalists’ and critics of RT, Stanley & Szabó (2000), and in fact criticised by Carston (2002).

<sup>26</sup>It may seem to some, both within and outside RT, that the implicature/explicature distinction provides for this sort of analysis (and parts of Carston 1998 suggest this). However, this would be in effect to equate explicature with the Gricean truth-conditional ‘what is said’, a notion at odds with the fundamental structure of RT (see Carston 2002, Chapter 2). The nature of explicature in the light of C&L’s views is worthy of discussion in its own right, but space does not permit here. See also footnote 6.

<sup>27</sup>See Wedgwood (2005, Chapter 5), Geurts (1998), Kadmon (2001) on the impact of factors like inherent scalarity, relationships to existential contexts and precision of quantificational value. See also Breheny (2005), who supports a kind of numeral/non-numeral distinction that is in some ways the opposite of Carston’s, with a more highly specified semantics for numerals.

trasts are demonstrably not categorical: it is easy enough to construct examples that show ‘truth-conditionality’ with non-numerals and ‘mere implicature’ with numerals (as I show in Wedgwood 2005, 163–167). As such, it is notable that the moderate contextualist element to Carston’s analysis—the concern that semantics should account directly for observations of Propositionality<sub>MC</sub>—is also the point that leads to empirical inaccuracy.

It is also notable that Horn (1992) is able to use Carston’s proposals to minimise damage to the neo-Gricean position: if there is something special about numerals that allows them alone to ‘affect truth conditions’ via inference, they can be sectioned off as an exceptional case, leaving the story otherwise unchanged<sup>28</sup>. Thus we have an illustration of how slippage into moderate contextualist reasoning is to the detriment of RT arguments both in terms of their internal coherence and their external impact.

It is important to note that this instance of covert moderate contextualism is (to my knowledge) an isolated blip in Carston’s work and that such things are, as they should be, rather rare in RT work as a whole. Nevertheless, it serves as a reminder of the importance of remaining suitably ‘radical’. This is significant even at the level of presentation. Moderate contextualist (in particular Gricean) assumptions are deeply ingrained in linguistic semantics and this may all too easily influence how linguists’ claims are understood, whether by other linguists or philosophers (as demonstrated by C&L’s misconceptions about RT). In this context, it is important for the relevance theorist to be quite clear about the implications of the pervasiveness of contextual influence on interpretation, and here it can be a real help to be armed with some sharp philosophical tools. In this way, C&L’s notions of the radical/moderate divide and the Mistaken Assumption can help relevance theorists to clarify just where certain dangers lie.

## 6 Conclusion

C&L’s criticisms of RT are misdirected, not least because RT has much more in common with their own ideas than they (and perhaps some relevance theorists) seem to realise. Both take as their starting point a rejection of the idea that grammatically encoded meaning should systematically account for ‘what is said’. Both recognise just two theoretically and psychologically important levels of meaning

---

<sup>28</sup>To explain the chronology: Horn draws on the unpublished Carston (1985), where the special, ‘sense-general’ treatment of numerals was first proposed, though this work does not include the formalisation using a ‘contextual variable’.

and both claim that these are: (1) that which is communicated by a given linguistic form independently of context and (2) that which is understood to be communicated in a particular context—which is an indeterminate number of propositions, among which implicatures and any sense of ‘what is said’ are derived by the same mechanisms.

Consequently, arguments from ‘content-sharing’ do not distinguish the two approaches, and as such it is good news for C&L that their attempt to dismiss RT on this basis proves to be unsustainable (as it confuses different notions of ‘content’). There remain just two identifiable differences. First, while C&L wash their hands of attempting to characterise the process of deriving Speech Act meaning from encoded meaning, RT proposes a set of inferential principles to do just this. Second, while RT makes no *a priori* claims for encoded meaning, C&L insist that minimal semantic content is propositional. This is in one sense a genuine difference between the two approaches and in another sense a rather spurious one. It causes C&L to include in minimal semantic content elements of meaning that would not enter into RT’s encoded meaning, and consequently clashes with their own requirement that minimal semantic content be shared across different contexts. However, it in any case plays no discernible role in the explanation of communication (supposed demonstrations of the importance of propositionality involving not minimal semantic content but what C&L call Speech Act content).

While this discussion of C&L and RT arguably reveals some unfortunate misunderstandings between the perspectives of philosophy of language and linguistic pragmatics (notably in C&L’s inappropriate application of an omniscient third party perspective to RT), it also suggests an encouraging convergence on ideas that represent an important challenge to dominant assumptions in semantic theory. From the perspective of philosophy, one important implication of this is that semantic minimalism and contextualism need not be viewed as incompatible—provided the semantics is minimal enough and the contextualism radical enough. Indeed, real semantic minimalists are in trouble without theories of pervasive context-dependent reasoning—something must fill in the gaps, which get bigger the more minimalist you are. Because of the implications for propositionality, this means that a thorough semantic minimalism cannot be ‘semantic’ in the sense that C&L and others might assume (Carston 2006)<sup>29</sup>. One thing I hope to have shown is that the significance of this should not be overplayed. In any case, it should not

---

<sup>29</sup>Korta & Perry (to appear) suggest how a slightly different notion of propositionality might be upheld that would not conflict with the need for minimal semantic content to be genuinely context-independent—but see also C&L’s (to appear *b*) response.

overwhelm the broader theoretical perspective, from which RT and C&L are very notably pushing in the same direction.

This is not only of philosophical interest; there are important implications for linguistic theory. Relevance theorists have tended to assume that RT can be used more or less as an adjunct to fairly conventional approaches to other parts of linguistic theory. It has been presented as a pragmatic framework to complement more or less independent work in syntax and semantics (see, for example, Carston 1988, Kempson 1988)—though of course there have always been some territorial claims made at the semantics-pragmatics border, as with any pragmatic framework. But in making the move away from the moderate contextualism of Gricean approaches, RT has more radical consequences, whether we like it or not. In effect, this constitutes a break from conventional perspectives on semantics, with both theoretical and methodological implications. As I have argued above, this is essentially the same break that C&L make. Neither they nor relevance theorists should be afraid to see through the implications of this<sup>30</sup>.

The logic of these approaches tells us that the nature of encoded meaning cannot be understood without active consideration of inferential contributions to meaning. This is nothing less than a reversal of conventional methodology, which tends to abstract away from inferential pragmatic processes as much as possible. This methodological reversal follows from the rejection of the compositional derivation of perceived truth conditions, upon which conventional semantic analysis rests. Once we reject this, as RT does in principle, the nature of encoded meaning becomes an entirely open question. And the consequences of this move go still further. A great deal of contemporary syntactic analysis is motivated by a wish to account for the perceived semantic character of a given sentence (hence the regular use of levels of syntactic representation like LF, *Logical Form*). In other words, a version of the Mistaken Assumption guides much syntactic theorising too. In rejecting this, both RT and C&L's approach imply a radical stance on the nature of syntax as well as semantics (at least in the context of currently popular approaches; see Wedgwood 2005 for extended discussion of this point). Because C&L offer no characterisation of the crucial inferential component in linguistic interpretation, they fail to provide a basis for following through the consequences of their own arguments in actual linguistic analysis. RT, on the other hand, provides

---

<sup>30</sup>Carston (2006) does hint that the consequences of RT for semantics are of a more profound kind than are generally recognised. C&L sometimes make out they are positively conservative in their 'semantic minimalism' (e.g. 2006a, 50), but elsewhere make clear that they are swimming against the tide: "Hardly any contemporary philosopher rejects [moderate contextualism]" (2005, 88).

a set of principles that can be the basis for exploring all these crucial issues.

## References

- Atlas, J. D. (2005), *Logic, Meaning and Conversation: Semantical Underdeterminacy, Implicature, and Their Interface*, Oxford University Press, Oxford and New York.
- Atlas, J. D. (to appear), How insensitive can you be? Meanings, propositions, context, and semantical underdeterminacy, in G. Preyer & G. Peter, eds, 'Context Sensitivity and Semantic Minimalism : New Essays on Semantics and Pragmatics', OUP, Oxford.
- Bezuidenhout, A. (2006), 'The coherence of contextualism', *Mind and Language* 21(1), 1–10.
- Blakemore, D. (1992), *Understanding Utterances: an Introduction to Pragmatics*, Blackwell, Oxford.
- Borg, E. (to appear), Minimalism versus contextualism in semantics, in G. Preyer & G. Peter, eds, 'Context Sensitivity and Semantic Minimalism : New Essays on Semantics and Pragmatics', OUP, Oxford.
- Breheeny, R. (2005), Some scalar implicatures really aren't quantity implicatures—but some's are, in E. Maier, C. Bary & J. Huitink, eds, 'Proceedings of SuB9', Nijmegen Centre for Semantics, Nijmegen, pp. 57–71.
- Cappelen, H. & Lepore, E. (2005), *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism*, Blackwell, Oxford.
- Cappelen, H. & Lepore, E. (2006a), 'Response', *Mind and Language* 21(1), 50–73.
- Cappelen, H. & Lepore, E. (2006b), Shared content, in E. Lepore & B. C. Smith, eds, 'Oxford Handbook of Philosophy of Language', Oxford University Press, Oxford.
- Cappelen, H. & Lepore, E. (to appear a), Relevance theory and shared content, in N. Burton-Roberts, ed., 'Advances in Pragmatics', Palgrave Macmillan, London.
- Cappelen, H. & Lepore, E. (to appear b), 'Reply to critics', *Philosophy and Phenomenological Research*.
- Carston, R. (1985), A reanalysis of some 'quantity implicatures'. Ms., University College London.
- Carston, R. (1988), Language and cognition, in F. Newmeyer, ed., 'Linguistics: the Cambridge Survey', Vol. III, Cambridge University Press, Cambridge.

- Carston, R. (1998), Informativeness, relevance and scalar implicature, in R. Carston & S. Uchida, eds, 'Relevance Theory: Applications and Implications', John Benjamins, Amsterdam, pp. 179–236.
- Carston, R. (2002), *Thoughts and Utterances: The Pragmatics of Explicit Communication*, Blackwell, Oxford.
- Carston, R. (2004), Explicature and semantics, in S. Davis & B. Gillon, eds, 'Semantics: A Reader', Oxford University Press, Oxford, pp. 817–845.
- Carston, R. (2006), 'Linguistic communication and the semantics/pragmatics distinction', *University College London Working Papers in Linguistics* **18**.
- Geurts, B. (1998), Scalars, in P. Ludewig & B. Geurts, eds, 'Lexicalische Semantik aus kognitiver Sicht', Gunter Narr, Tübingen, pp. 95–118.
- Grice, H. P. (1975), Logic and conversation, in P. Cole & J. Morgan, eds, 'Syntax and Semantics 3: Speech Acts', Academic Press, New York, pp. 41–58. Reprinted in Grice (1989): 22–40.
- Grice, H. P. (1989), *Studies in the Way of Words*, Harvard University Press, Cambridge, Mass.
- Gross, S. (2006), 'Can one sincerely say what one doesn't believe?', *Mind and Language* **21**(1), 11–20.
- Horn, L. (1972), On the semantic properties of logical operators in English. Ph.D. thesis, UCLA. Distributed by Indiana University Linguistics Club, 1976.
- Horn, L. (1992), The said and the unsaid, in 'Ohio State University Working Papers in Linguistics 40 (Proceedings of SALT II)', pp. 163–192.
- Kadmon, N. (2001), *Formal Pragmatics: Semantics, Pragmatics, Presupposition, and Focus*, Blackwell, Malden, Mass. and Oxford.
- Kempson, R. (1988), Introduction, in R. Kempson, ed., 'Mental Representations: The Interface between Language and Reality', Cambridge University Press, Cambridge.
- Korta, K. & Perry, J. (to appear), 'Varieties of minimalist semantics', *Philosophy and Phenomenological Research*.
- Larson, R. & Segal, G. (1995), *Knowledge of Meaning: An Introduction to Semantic Theory*, M.I.T. Press, Cambridge, Mass.
- MacKenzie, I. (2002), *Paradigms of Reading: Relevance Theory and Deconstruction*, Palgrave Macmillan, Basingstoke.
- Montminy, M. (2006), 'Semantic content, truth conditions and context', *Linguistics and Philosophy* **29**(1), 1–26.
- Sperber, D. & Wilson, D. (1986), *Relevance: Communication and Cognition*, Blackwell, Oxford. (Second Edition: 1995).
- Stanley, J. & Szabó, Z. G. (2000), 'On quantifier domain restriction', *Mind and*

- Language* **15**(2–3), 219–261.
- Travis, C. (2006), ‘Insensitive semantics’, *Mind and Language* **21**(1), 39–49.
- Wedgwood, D. (2005), *Shifting the Focus: From Static Structures to the Dynamics of Interpretation*, Elsevier, Oxford.
- Wilson, D. & Sperber, D. (2004), Relevance Theory, in L. Horn & G. Ward, eds, ‘The Handbook of Pragmatics’, Blackwell, Oxford, pp. 607–632.