

# Shared assumptions: semantic minimalism and Relevance Theory

Daniel Wedgwood  
School of Philosophy, Psychology and Language Sciences,  
University of Edinburgh

June 8, 2006

DRAFT

## Abstract

Cappelen & Lepore (2006a) claim that linguistic communication requires *shared content* and criticise Relevance Theory on the grounds that it makes content sharing impossible. I argue that this criticism rests upon two important errors. The first is a flawed understanding of Relevance Theory, which causes Cappelen & Lepore to apply the perspective of an objective, omniscient third party to parts of Relevance Theory that in fact only make reference to (and only depend upon) the subjective judgements made by the addressee of an utterance. The second is a confusion about different definitions of *content*. The kind of evidence that Cappelen & Lepore produce in fact involves the communication of what they term *Speech Act content*, a kind of content that few if any approaches require to be perfectly 'shared'. Certainly, Cappelen & Lepore's own (2005) position makes no such requirement. Consequently, their arguments have no more force against Relevance Theory than they would if directed at their own ideas. Having shown this, I compare Relevance Theory with Cappelen & Lepore's *semantic minimalism* more generally and conclude that most of the apparent differences between them lie in terminology and rhetorical style. Instead of highlighting these spurious differences, practitioners of the two approaches have a good deal to learn from each other. In particular, I argue that relevance theorists should avoid lapsing into a brand of semantic orthodoxy that Cappelen & Lepore call *moderate contextualism*, something that requires them to emulate Cappelen & Lepore's uncompromising minimalism with regard to encoded meaning.

# 1 Overview

At the beginning of their (2006a) 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.

Unfortunately, for a notion that is apparently so fundamental, C&L leave a good deal of confusion surrounding the definition of *content*. This begins with the mixture of ideas in the above quotation. The first sentence states that to successfully perform a speech act (an assertion, claim, conjecture, etc.) is to share content. This a *process* of sharing. The second sentence refers not to speech acts but to thoughts and (therefore) not to a process but a *state* in which something could be said to be shared. The third sentence introduces the idea of evaluating what is conveyed by the speech acts of another person; a second kind of process and one that does not obviously involve sharing in any sense. Consequently, the above passage fails make clear what C&L really mean when they continue the above passage with following claim, which is the central argument of their paper: “Relevance Theory. however, implies that content sharing is impossible”.

Do C&L wish to claim that Relevance Theory (henceforth RT) implies that performing speech acts is impossible? Or do they rather mean that RT implies that people cannot entertain the same thoughts as each other? Or that RT implies the impossibility of evaluating whether what another person says is “true, false, interesting, obscene, original or offensive” (etc.)? If RT or any comparable theory were to imply any of the above, it would be both surprising and crushing for that theory, as these kinds of activities are in effect the primary data that such a theory is intended to explain. Below I argue not only that RT has no intrinsic problem accounting for such phenomena but also that C&L’s own approach, as set out in C&L (2005), works just like RT with regard to these phenomena.

In any case, it is not clear to me that C&L’s article even makes any of the above claims as such. Picking through their uses of the term ‘(shared) content’, it is generally difficult to tell exactly what RT is being accused of. 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 discuss the latter in this context. 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 kind of content—insofar as it is made explicit anywhere—is shown to be equivalent to that of RT in all crucial ways. 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. They and many relevance theorists seem convinced that they differ, but this is mainly on the basis of different rhetorical strategies. These in turn stem from different theoretical backgrounds and different kinds of reaction to the dominant paradigm in linguistic semantics, which C&L (2005) call *Moderate Contextualism*. I argue below that a number of C&L's criticisms of RT in fact have force only as arguments against moderate contextualism. On the other hand, I also show that some avowed relevance theorists have been guilty of bringing moderate contextualist assumptions and practices into analyses supposedly made within RT. As such, relevance theorists can learn from C&L's rigorous, if sometimes misdirected, views on contextualism.

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 call into question (without attempting a full answer) whether this difference is as substantial as it may sound—at any rate it does not affect the arguments surrounding C&L's criticisms of RT. C&L state the minimal semantic content of a given utterance in the language of Davidsonian truth-conditional semantics, but what they actually claim thereby amounts to the following, rather less than contentious, statement: the linguistic form of each utterance—i.e. each sentence—conveys *something* that is consistent across all contexts and that all speakers of a given dialect can share by producing that utterance. One thing C&L never do is venture to suggest what exactly this *something* is, beyond saying that it is whatever makes the sentence true (for them, to go

any further is the job of metaphysicians, not semanticists). Unlike many critics of Cappelen & Lepore (2005), I find the rights and wrongs of invoking truth conditions in this way to be somewhat beside the point (as I indicate below, C&L may have a point here, but they also have significant problems). More importantly, this view of minimal semantic content leaves C&L relying on exactly the kind of processes that RT aims to characterise, without themselves venturing such a characterisation. As such, RT is not only just as theoretically coherent as C&L's position, but also considerably more useful.

C&L are aware of the accusation that they are 'disguised Relevance Theorists' (2005, 181). In section 6, I briefly discuss their most explicit refutation of this accusation and show that it depends on drawing unsupported implications from the one clearly identifiable difference between their approach and RT. I conclude that C&L and relevance theorists have more to learn from each other than they have to gain from attacking one another.

## 2 Relevance Theory and shared content

### 2.1 Inferring and sharing

Sperber & Wilson's (1986) Relevance Theory 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, to my mind, the most coherent radical contextualist framework (as I argue at length in Wedgwood 2005, Chapter 2), in part 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 utter-

ance 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.

Without wanting to leap naïvely into non-linguistic areas of philosophy, I think it is worth placing this idea in a broader context (which will help also to provide the basis for my defence of RT in the face of C&L's criticisms). I take the following to be extremely basic, uncontroversial (at the very least, patently sensible) observations about the general limitations on human beings' processing of all kinds of information. Without in any way compromising a realist philosophy, we have to accept the inherent subjectivity of our individual experiences and beliefs: access to the rest of the universe is inevitably indirect, being mediated first by our physical senses and then by our abilities to process and reason over the data that we receive from them. These senses and mental abilities have evolved to reflect the outside universe in particular, skewed ways and both are fallible even in these terms. It follows that anything that a given individual considers to be a piece of knowledge (or a belief, a suspicion, etc.) is something like a scientist's theories (and working hypotheses, etc.): best guesses that she has more or less reason to believe in and which may to a greater or lesser extent provide a safe basis for drawing further conclusions. Everyday experience tells us that what different people take to be true, even unquestionably true, varies enormously; on the other hand, there are certain things that pretty much everyone agrees upon (either because there is just so much evidence for them or because we are innately set up to find certain things to be the case).

The point of this digression is twofold. First, I want to emphasise that RT is merely part of a sensible, even necessary response to the bigger picture. RT relies on the notion that interlocutors infer each other's intentions. This sometimes makes semanticists balk; for some, it makes RT a non-starter as a linguistic theory (or even as a theory that directly impacts on linguistic competence; e.g. Asher & Lascarides 2003, 76). But ultimately a theory dealing with communication and cognition will have to mesh with the bigger picture—and, conveniently or otherwise, this is a picture in which people spend a lot of their time making educated

guesses about the way things are, based on the best evidence (as they see it) before proceeding to accept a number of these as fact. Communication is closely tied to the formation and expression of (what an individual takes to be) knowledge, so if this is the fundamental nature of what we ‘know’, we might as well build it into our models of communication from the beginning (I’ll come to RT’s particular take on ‘what people take to be knowledge’ below). Furthermore, if we must assume this kind of hypothesis-forming, best-guessing behaviour to explain people’s most basic understanding of their world, then we can take certain abilities for granted in our explanations of communication—and we might as well take what comes for free<sup>1</sup>.

Second, the bigger picture is directly related to the reasons why RT does not, as C&L claim, make shared content impossible—at least not in any philosophically troubling way. As I outline below, the basis of C&L’s (2006a) argument could only be that RT’s ‘radical contextualist’ denial of directly communicated propositions leads to the impossibility of truly sharing content, on the grounds that it leaves one relying on interlocutors’ best guesses at what the speaker of an utterance meant to convey. This means (the argument goes) that RT can’t make reference to sharing identical content—and C&L go on to argue convincingly that mere *similarity* of content is not a suitable basis for explaining communication. What the bigger picture tells us is that people have to rely on best guesses (or best corroborated hypotheses) in a much wider context than just communication, and that we have effective strategies for doing so—if far from infallibly, then well enough to get by. As in other domains of knowledge-acquisition, people don’t need *objective* identity of content in order to communicate successfully (enough); they just need reason enough to believe that what they have recovered from a communicative act is identical to what the speaker intended. RT is 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

---

<sup>1</sup>‘Radical contextualists’, when promoting a pragmatic analysis of some linguistic data, often point out that some form of intention-inferring abilities are independently necessary, on the basis of such context-dependent linguistic phenomena as conversational implicature, sarcasm, disambiguation and reference assignment. I want to emphasise here that these linguistic phenomena are mere manifestations of the relevant inferential abilities, not the primary motivation for positing them.

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<sup>2</sup>.

## 2.2 The structure of Relevance Theory and semantic minimalism

In this section, I offer a skeletal outline of how RT works (for proper introductions to the framework, see Sperber & Wilson 1986, Blakemore 1992, Carston 2002, Wilson & Sperber 2004). My aims here are (a) to provide the basis for comparison to C&L's approach, below, (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 connotations 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 do regularly communicate with each other, meaning 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: they argue that “what speakers say (assert, claim, state, etc.)” is a matter of *Speech Act content*, rather than semantics (2005, 13) and that Speech Act content is heavily context-dependent (2005, Chapter 13). Meanwhile, their emphasis on the ‘sharing of content’ shows their commitment to (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 the generation of cancellable (i.e. ‘non-truth-conditional’) 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 much a more basic position

---

<sup>2</sup>One possible question, which I do not pursue in the present article, is whether C&L's position is the result of evaluating RT from a straightforwardly incompatible philosophical viewpoint. Cappelen & Lepore (2005, 2006a) do not make explicit reference to the foundational perspective on semantics provided by the likes of Fodor (1987), but this seems to me to display the same kind of dependence upon an objective notion of truth and knowledge that C&L use to judge RT. I would suggest that RT more naturally fits within the kind of mentalist approach to the foundations of semantics promoted by Jackendoff (2002, Chapter 10) (though, as more generally, I do not claim to speak for all confessed relevance theorists in this respect).

on what kinds of meaning (or ‘content’) there must be: some meaning is encoded in linguistic forms and some is inferred. The meaning encoded in a given linguistic expression 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 is of course the encoded meaning of the linguistic expression used. But there must also be a means of evaluating this, in the light of some kind of further evidence, in order to understand its particular significance on a given occasion. Such further evidence is extra-linguistic and may in principle be drawn from any information that is available to the individual. But (in line with the bigger picture, as set out above), ‘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. it is manifest to each person that it is manifest to the other. 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, 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.

This gives us a handle on how language users might calculate which assump-

tions to bring to bear on the interpretation of a given utterance. Other things being equal, only the most easily accessible assumptions should be involved. What might make other things unequal? From the speaker's point of view, this might be the desire to communicate something in particular which inherently requires reference to something less easily accessible. From the addressee's perspective, however, extra processing effort would only be justified by the promise of some communicative 'reward' commensurate to that effort. Generalising over these two perspectives, we can say that a more obscure, more costly to access part of the mutual cognitive environment should only be invoked if it is truly relevant.

RT is based upon the idea that addressees can in this way use the effort required to process an utterance to guide them to the intended interpretation of it. If a potential interpretation does not seem to constitute a worthwhile piece of communication (e.g. it mostly tells the addressee what she already knows or tells her things that don't seem to relate to what she knows), then this is an indication that it should be worth working a little harder to find another interpretation, using less obvious manifest assumptions and/or making some extra deductive steps. On the other hand, if a worthwhile interpretation does come up on the basis of easily accessible mutually manifest assumptions, then the addressee may reasonably assume that this was the intended interpretation of the utterance.

More technically, RT characterises this process in terms of a 'presumption of optimal relevance', where relevance is conceived of as a ratio of communicative rewards to processing effort. These 'rewards' are of a very abstract kind: they are what RT calls *cognitive effects*: strengthening or weakening/eliminating existing assumptions or providing new assumptions that interact with existing assumptions to allow for the deduction of further assumptions<sup>3</sup>.

The 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

---

<sup>3</sup>The apparently rather complicated formulation of the last kind of cognitive effect merely reflects the intuitive fact that new assumptions with no connection to existing ones are not relevant to an individual.

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: encoded meaning and inferred meaning. Those familiar with the framework will no doubt have heard of other apparently significant distinctions from the RT literature, such as implicature versus explicature. As I discuss below, in section 5.1, such distinctions may be treated as key theoretical notions by some authors, but are not ultimately more than expository devices according to the underlying logic of 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>4</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.

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 cognitively, psychologically and communicatively significant levels of meaning: what they call minimal semantic content 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:

[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 5.1). It may be that relevance theorists tend to express

---

<sup>4</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.

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. On the other hand, the basis of C&L’s claim here is a conception of truth conditions which is so general and so far divorced from usual notions of truth-conditionality that it is unclear just how much substance it should be credited with. Note that I am not claiming here that C&L’s particular notion of propositionality is necessarily illegitimate; for me the jury is still out on this. What I am claiming is that the nature of C&L’s propositionality makes this one notable difference between their position and RT much less significant than it might appear.

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 next section.

### 3 C&L’s criticisms of RT

#### 3.1 ‘Non-shared content’ in reality and in theory

C&L claim that RT, like other radical contextualist approaches, contains what they call a *non-shared content principle* (or *NSC*):

(NSC): When a speaker utters a sentence, *S*, thereby intending to communicate the proposition that *p*, the audience will not grasp *p*. Instead, she will interpret the speaker to have intended to communicate some proposition (or set of propositions) R-related to *p*. (Cappelen & Lepore 2006a, 3\*\*\*)

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)

C&L claim that the ‘R-relation’ referred to in NSC must be *similarity*. 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 formulation of NSC, as a characterisation of RT. First, it represents a particularly strong statement that is not clearly entailed by RT as a whole, or indeed by the above quotations. What RT recognises is that the addressee *may* not grasp an intended propositional meaning *p*; not 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 actually do so (or come negligibly close to it) regularly they would probably soon find reasons *not* to believe any such thing. However, the precise degree of reliability attributed to linguistic communication is not crucial to the present argument. As I show below, 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 is ultimately affected by the truth or otherwise of this idea, not by whether they have overstated NSC as a characterisation of RT.

The second, more important problem with C&L’s NSC is that it implies that the ‘R-relation’ of similarity is an active component of RT. It is not. C&L’s NSC describes what may well be the case in a given act of communication 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 perceptions of what must be the intended meaning of an utterance (and how these perceptions come to be formed). It is this subjective assessment of intended meaning that drives RT as an explanatory framework and this, I argue, is also all that is required to understand the kinds of communication that C&L discuss 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 third party, without suggesting that

this is part of the explanation of how communication is achieved. This comes through once the quotation is placed in the context of this 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. Put more bluntly, 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 I have some sympathy with them: 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 wholly successful 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. Nevertheless, the process that 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 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.

Drawing the distinction between similarity as an objectively identifiable fact and similarity as a (specious) component of RT is thus an important part of my reply to C&L. However, the above discussion does not in itself address the substance of C&L’s arguments, which are principally built on how the putative principle NSC interacts with certain linguistic facts (notably the interpretation of reported speech). The next step in my argument is therefore to reconsider these interactions, bearing in mind that NSC is an objective exterior description of (some cases of) communication, to be distinguished from the internal, subjective processes by which RT claims to explain communication.

### **3.2 C&L’s empirical arguments**

The structure of C&L’s (2006a) argument is as follows. They present a number of uses of language that supposedly necessitate the existence of ‘shared content’; let us call these C&L’s primary test cases. These are in effect linguistic expressions of judgements made about other linguistic utterances and as such they provide a way for C&L to use judgements of truth or falsity to get at the meaning conveyed by an utterance (I suggest below that this is in itself problematic, but this is not the present concern). They proceed to argue that a relevance theorist could only

accommodate such phenomena by adopting one of two strategies, which they call the *Revisionist Strategy* and the *Conservative Strategy*. They then produce further illustrations of linguistic behaviour which they believe show each strategy to be empirically unsustainable and/or in some way deeply philosophically problematic. I believe that only the first of these strategies would ever be adopted by a (coherent) relevance theorist and that, far from being revisionist, this strategy is in fact clearly implied by the basic nature of RT (in ways already discussed above). None of the further illustrations of linguistic behaviour that C&L proffer actually conflict with this aspect of RT, once one properly distinguishes the roles of subjective and objective judgement in RT explanations.

I first deal with one of C&L's (2006a) primary test cases, which straightforwardly illustrates the argument of the previous section and as such can be dealt with quickly. Following this I mention C&L's other primary test cases, which lead to discussion of the so-called Revisionist and Conservative strategies.

### 3.2.1 Assessments of assertions

C&L argue that 'shared content' is required to make sense of our ability to assess the linguistically conveyed assertions of others. Their examples involve an original utterance of some sentence *S* by a speaker named Naomi and the following kind of responses to it: *That's true/questionable/unacceptable/clever*, or *I agree with Naomi*. C&L's argument here is that "the speaker [of such responses] doesn't mean to evaluate a proposition similar to the one asserted by Naomi's utterance; the speaker [...] means to evaluate the very proposition/thought that Naomi expressed." (Cappelen & Lepore 2006a, 7\*\*\*).

This is quite correct and entirely unproblematic for RT. The speaker who responds to Naomi by uttering *That's true* clearly believes that he has understood what Naomi meant by her original utterance. Whether this is objectively the case or not makes no difference to our ability to understand what is going on in the dialogue: the second speaker is concurring with what he believes Naomi to have intentionally communicated. There is no sense in which he *means* merely "to evaluate a proposition similar to the one asserted by Naomi's utterance", whether or not he inadvertently does so (and note that this is the case even if he is conscious of the possibility of having inadvertently done so). C&L's objection therefore misses its target in a way that illustrates the point made in section 3.1: their mistake here is plainly one of treating NSC as part of the mechanisms of RT, rather than as a possible consequence of RT when viewed from an omniscient third party perspective.

### 3.2.2 Speech reports and belief attributions

Another of C&L's other primary test cases is also one of their favourite talking points in their (2005) book: indirect speech reports. They ask us to imagine again that Naomi has uttered a sentence *S*, but this time 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>5</sup>. C&L then observe that "If NSC is correct, the proposition attributed to Naomi will (typically) be similar to, but not identical to, the proposition that Naomi intended to communicate with her utterance" (2006a, 5\*\*\*). As discussed in section 3.1, a more appropriate version of NSC would necessitate replacing *will (typically) be similar* with *may be similar* in this quotation, but C&L's subsequent remarks are more important: "A proponent of NSC [...] 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.

Very similar reasoning is applied to another primary test case: the fact that we often attribute beliefs to our interlocutors based on what they say. Imagine that because of Naomi's utterance of *S*, Donald concludes that *Naomi believes that p*. Again, by NSC, *p* is (or, better, may be) only similar to the proposition that Naomi intended to convey. C&L conclude that the committed relevance theorist must sustain either that (a) the very common process of taking people's sincere assertions to reflect their beliefs actually yields (I'd say: may yield) false attributions of belief or (b) we can judge belief attributions to be true even when they describe beliefs that are only similar to the ones actually held by the person they are attributed to. Again, (a) and (b) are said to be manifestations of the Revisionist and Conservative strategies.

### 3.2.3 'Conservatism'

I take it that no proponent of RT would consider adopting the so-called Conservative strategy, which amounts to re-defining our very notions of truth and identity such that 'similar' propositions could be identified as being truth-conditionally

---

<sup>5</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.

equivalent. As C&L rightly point out, this would wreak havoc with all kinds of inferential procedures. They also point out 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 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. 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 and the belief attribution are true as such iff Donald's judgement is right: iff  $p$  is (objectively) identical to Naomi's intended assertion. Otherwise, they will be, strictly speaking, false. But either way he can carry out the report or the attribution sincerely and effectively. Ultimately, whether the truth or falsity of his statements 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 understand one another's beliefs 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 attributed her with a belief 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.

#### **3.2.4 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 henceforth refer to this as the R-strategy (the *R* can then be taken to stand 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.

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>6</sup>. This 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 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>7</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 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 (as in their NSC), which implies that they believe that content sharing (in the sense of reproducing communicators' thoughts) need only be possible in order to facilitate communication. 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 (not 'hopelessly fallible' or even

---

<sup>6</sup>At least, this is the line they push in their 2006*a* article. On the other hand, see C&L (2005, 206) for apparent recognition of the more problematic aspects of verbal communication.

<sup>7</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.

‘inevitably inaccurate’<sup>8</sup>). 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, I think both of these things are the case. C&L really do seem have missed the idea that the Principle of Relevance is meant principally to explain how addressees do regularly 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” (2006a, 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 RT does this, but one cannot criticise RT as though it fails to address this problem, when that is its very *raison d’être*. The fundamental argument of RT has always been that recognition of speakers’ communicative intentions would indeed be a minor miracle in the absence of some set of guiding principles like those of RT: principles which enable interlocutors to assess just how they can bring appropriate assumptions to bear on the interpretation of linguistic expressions or, as speakers, use linguistic expressions to trigger appropriate interpretations. 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 this.

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, I suggested in section 2.1 that RT’s recognition of a

---

<sup>8</sup>In this context, it is worth noting that RT has itself 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’ (intellectual company that most serious philosophers and scientists would not wish to keep, I venture to suggest) 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.

lack of guarantees merely puts communication into the context of very general facts about human processing of data about the outside world: little if anything is truly guaranteed, but 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 does not in principle contradict the idea that the intended propositional content of utterances is not conveyed (hence shared) directly and infallibly.

Certainly, 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. As I have indicated above, the latter observations are not only consistent with RT, but are precisely what RT sets out to explain.

Nevertheless, there is clearly a danger here of to some extent merely trading intuitions: some of us are struck by the fallibility of communication, others by just how much we can achieve by it. We can go beyond this by looking in more detail at what C&L’s position implies; at the nature of the relevant kind of ‘content’ itself. This shows C&L’s position to be flawed in the following way: *any* account of the kind of content sharing that is pertinent to C&L’s observations must rely on indirect calculations of intended meaning. This includes C&L’s own analytical standpoint, as I discuss in more detail 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 kind of content that C&L argue to be infallibly shared—indeed, the only kind of content that *could* be, on any view—is that of linguistically encoded meaning
- but 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 point, that C&L in practice argue for the sharing of *linguistically encoded* content, is implicit in their criticisms of RT and explicit in their extended (2005) 'defence of semantic minimalism'. That this is the only kind of content that could even in principle be shared may also be deduced from a consideration of the only logical alternative: that Speech Act content could be directly, infallibly shared. As pointed out above, this is something that no-one, C&L included, would want to maintain. Whatever the nature of encoded meaning may be, it is uncontroversial that much Speech Act content is only indirectly related to it and is highly context-dependent. As such, it could never be claimed that the very act of producing some utterance causes such content to be shared. I return below to the nature of C&L's own view of Speech Act content and its relation to RT ideas.

#### 4.1 Minimal semantic content

Linguistically encoded meaning of course *has* to be directly shared. This is what it means to be directly encoded meaning: it is whatever two speakers of the same dialect can recover directly from a given piece of linguistic material. Where RT and C&L differ is therefore not on the existence of directly encoded (and hence directly shared) meaning, but on its nature. Specifically, C&L claim that all encoded meaning is propositional. This is their 'minimal semantic content'.

It is important in this context to appreciate just how minimal C&L's minimal semantic content is. As we have seen, when they claim that encoded meaning is propositional, they do not mean it in the sense of the far more widespread point of view that they call moderate contextualism. That is, they explicitly reject the idea that intuitively derived truth conditions should be taken to correspond to the meaning encoded in a sentence (even *modulo* saturation of indexicals). Instead, the kind of minimal propositions that C&L have in mind concern *just whatever makes a sentence true* across all contexts. This may be something that we have little or no access to either as speakers or analysts—C&L argue that specifying such truth conditions is the business of metaphysics, not semantics. This leaves them stating only meanings that strike most analysts as being quite circular, such as (1):

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

Typically, this strategy attracts criticism along the lines that (1) does not really contain truth conditions. It is important to emphasise that this is not the nature of my argument with C&L. Though such criticisms may come from those who might happily accept C&L’s label *radical contextualist*, I do not believe that they are appropriate from such a viewpoint. I am convinced by C&L’s (2005, Chapter 11) retort that this kind of criticism falls wide of the mark, because of the distinction between semantics and metaphysics (as they put it), and represents a confusion of intrinsic meaning with merely intuitively graspable meaning. Still, it seems to me that there is one crucial matter that they fail to make explicit with regard to this argument. This is the fact that (presumably) their statements of truth conditions are to be understood, as is standard, *modulo* some denotation assignment function. Once this is considered, the fundamentals of formal semantics do seem to presuppose something like C&L’s position rather than the moderate contextualist position that is regularly adopted in practice in semantic analysis.

Assignment functions are rarely discussed, regularly being simply assumed and suppressed in semanticists’ representations for ease on the eye. Nevertheless, they play a crucial role, one that has been identified with the effects of context (e.g. Heim & Kratzer 1998, 242). In other words, assignment functions are in a sense promissory notes for a theory of pragmatics. Now if, as is usual, one directly translates words like *table* and *love* into predicates like *table’* and *love’*, one implicitly apportions to assignment functions (hence a theory of pragmatics) *all* the work that lies between those words and any intuitively graspable propositional meaning. No conventional semanticist feels that *table’* and *love’* fail to contribute adequately to the composition of truth conditions, despite the fact that what they contribute could vary massively with different assignment functions. The reason this is not seen to be problematic is precisely because the role of the assignment function is widely recognised here; everyone is comfortable simply assuming that values are relativised to assignments at this level. Now, just the same applies in principle to more complex expressions: there is a gap between a sentence meaning (the meanings of words put together in a certain way) and intuitively accessible truth conditions that the formal semanticist should be prepared to leave to assignment functions—meaning ultimately to a theory of pragmatics. Precisely what these functions consist in is universally recognised to be quite beyond semantics. But this means that the precise extent of their influence cannot

be known either. From this it follows in turn that the precise relationship between encoded meaning and intuitively accessible truth conditions is not knowable from the semanticist's perspective, only via active consideration of pragmatic theory. A semanticist should therefore recognise that the truth conditions of a propositional meaning that is stated *modulo* assignment functions may be unknowable. This brings us back to what is essentially C&L's position.

I do not therefore object to C&L's notion of (or expression of) minimal semantic content in principle (though I would in any case be tempted to call it less than entirely helpful). This does not mean that one must accept the propositionality of encoded meanings, however. From my point of view (as in principle for C&L), the question of whether all linguistically encoded meanings can be legitimately treated as in (1) is an empirical one. I address this issue below, in section 5.

For now, the issue is content-sharing. Does the propositionality or otherwise of minimal semantic content make a difference to this? C&L clearly assume that it does, but no part of their argumentation on the role of minimal semantic content in communication actually seems to depend on its propositional status. Consider the following quotation:

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 set of linguistic forms (such as a sentence). It does not depend upon this constant meaning having any particular properties, such as being propositional. Thus, 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 (see further section 6). Rather more importantly where C&L's (2006a) arguments are concerned, the role of minimal semantic content simply does not extend to the functions of shared content that are set out in C&L's criticisms of RT. These must involve a different kind of content altogether.

## 4.2 The content that is shared

Recall the kinds of observations that C&L claim make shared content indispensable: co-ordinated action, collective deliberation, linguistic communication justi-

fyng 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 point out that *I just said steel is strong enough; I never said strong enough to support the roof*. At least where I come from, *reasonable* isn't the word one would use to describe this kind of utterance. 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 2006a, 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 one of C&L's primary test cases, the assessment of assertions. The SA 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 (2)<sup>9</sup>:

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

---

<sup>9</sup>The parallel with C&L's examples is not perfect, because assessments of truth cannot be made about implicatures, which are by definition cancellable. This does not affect the point here, which is to demonstrate that in general the ability to assess utterances relies on the recovery of SA content, not minimal semantic content.

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 SA content. Therefore, what is necessarily shared 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>10</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, not even the speaker. (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).

What C&L claim with regard to SA meaning is in fact nothing more than that, like implicatures, it must be calculable in context on the basis of the meaning directly encoded in language. As such, C&L make no more claims for SA meaning than relevance theorists do—certainly, there is no guarantee that SA content will be shared perfectly. Consequently, C&L’s arguments regarding ‘shared content’ do not address any point on which they and RT actually differ. It follows that collective deliberation, inferences of belief, and so on, are 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.

One way in which the two approaches do differ is that RT provides an explanation for how SA communication is achieved as often as it is, and as such

---

<sup>10</sup>On the other hand, the following quotation also shows a point of C&L divergence between C&L and RT, on the issue of what C&L call Original Utterance Centrism. I argue below that this is a further example of C&L’s failure to separate subjective and objective notions of meaning/content conveyed. See section 6.

provides a basis for the explanation all the things that C&L take to necessitate ‘shared content’. Note that RT’s pursuit of a ‘reliable but fallible’ strategy for the recovery of intended meaning—just what C&L criticise RT for—is advantageous in this regard. Since no-one suggests that SA meaning is communicated infallibly, an RT-style strategy (in this sense) is exactly what is required to explain the phenomena that C&L discuss.

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 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 speech act content is most likely quite impossible, one can theorise about the principles that people employ in their attempts to derive such content. Furthermore, such principles must concern how to infer speaker’s communicative intentions.

To a considerable extent, this is just what the traditional notion of a Speech Act consists in (although SA theory might not commonly be spoken of in these terms). To borrow from the title of Austin’s (1962) seminal work, to understand SA content is to work out what the speaker is *doing with words*—and, where ostensive communication is concerned, this is inseparable from the speaker’s intentions. As such, there is no avoiding an inferential, intention-based approach at this level of analysis. Moreover, C&L’s ‘semantic minimalism’, which, like RT, recognises that encoded meanings (or minimal semantic content) may not be even a salient part of SA meaning, is more clearly reliant on a theory of intention-based inference that most approaches.

To summarise this section, C&L’s arguments from shared content do not speak against RT, but rather demonstrate the need for an approach broadly like RT, because these arguments logically must refer to SA meaning, not to C&L’s minimal semantic content. C&L’s own position presupposes some indirect means of establishing SA meaning (as opposed to direct sharing of it), which ultimately means that they presuppose something that has just those features RT that they try to criticise.

The question then arises to what extent C&L’s approach really differs from RT, whether the ostensible differences are substantial and whether either side of any substantive difference can be shown to be theoretically or empirically better

supported. This is the business of the following section.

## 5 Tests for minimal semantic content

As argued above, there is just one clear and notable difference between the claims of RT and those of C&L, once terminological differences are accounted for. This is C&L's insistence that minimal semantic content (encoded meaning) is propositional. For relevance theorists, encoded meaning is typically, perhaps always (Carston 2002), underdetermined with respect to propositional meaning. As I argued above, these claims cannot necessarily be directly contrasted, since C&L's notion of propositionality is not the same one that relevance theorists are usually keen to dispel. However, C&L's line on the propositionality of minimal semantic content remains something that does not exist in RT, so it is worth considering the basis of this as well as the extent to which it constitutes a significant difference for any practical purposes.

C&L do provide a putative empirical basis to their circular-looking truth conditions (as they surely must, if these are to have any real substance). 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 intuited propositions and they supposedly constitute C&L's most substantial objection to radical contextualism<sup>11</sup>—and they also provide the only real evidence of propositionality in minimal semantic content. These tests are close relations of the (2006a) primary test cases 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. 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, (3b) would be false as a report of (3a) because of the context-sensitivity of the pronoun *I* (Cappelen & Lepore 2005, 89). 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.

---

<sup>11</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.

- (3) 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.’

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 moderate contextualist who would deal with the ‘incomplete truth conditions’ of strings like *John is ready* by positing covert indexical elements, but not necessarily for 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 context in other ways<sup>12</sup>. Again, the use of assignment functions in formal semantics reflects this viewpoint (despite the common philosophical context in which such formalism is typically employed): indexicals are certainly significantly different to other lexical items in formal representations, but this does not mean that other items are entirely free of the influence of context.

Up to this point, then, I do not see C&L’s viewpoint diverging in principle from radical contextualism. Their argument regarding true indexicals and context-sensitivity only has force as an objection to moderate contextualism, even though they present it as an objection to the radical kind. As illustrated in section 5.2, it may be that many professed radical contextualists in practice behave more like moderate contextualists in positing unnecessary material (covert indexicals and the like) to account for context sensitivity, but radical contextualism in theory has no problem accommodating C&L’s observations.

If these tests do not undermine radical contextualism (hence RT) as such, they might still demonstrate the key difference between C&L and RT: the 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

---

<sup>12</sup>See Atlas (2006) for a careful elaboration of the point that recognising a case of context-sensitivity does not entail a claim of indexicality of any kind.

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 truthfully report that *In both C1 and C2, Nina said that John is ready*. C&L take this to be an argument against radical contextualism because this view would require the impossible situation that the complement of *said* expresses two distinct propositions simultaneously.

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—and then it is simply not a test for propositionality<sup>13</sup>.

There are various ways in which indirect speech reports of the kind that C&L rely on can fail to report a single proposition. In this respect, the evidence suggests that these speech reports are not necessarily as distinct from direct reporting as C&L's reasoning requires them to be. 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 (4):

- (4) 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

---

<sup>13</sup>In this context 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 (2006b). A grammatical context whose significance is inconsistent is not a test for anything.

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.

The problems are not restricted to ambiguities. C&L's test fails to produce a clear judgement of truth or falsity in just those situations that might make them really useful to linguists: when we need to determine whether grammar or pragmatic principles (or some particular combination of the two) should account for some phenomenon of meaning. An example of such a phenomenon is 'association with focus', as expressed by a certain kind of pitch accent in English<sup>14</sup>. Imagine that in C1 Bert utters (5a). Then in C2 he utters (5b).

- (5)       a.    Jim only introduced BOB to Alice.  
          b.    Jim only introduced Bob to ALICE.

The semantic orthodoxy with such cases is that the location of the main pitch accent induces a restructuring of the semantic form of the sentence such that the domain of the quantificational operator *only* can be fixed in two different ways. On the other hand, one might want to argue that the compositionally derived part of the meaning—the minimal semantic content—is determined by the string of words alone and that prosody merely provides a trigger for inferring the intuitive reading of each utterance. If C&L's indirect report test truly identifies propositional minimal semantic content, it should be of some use in resolving this debate. Can I then truthfully report that *In C1 and C2, Bert said that Jim only introduced Bob to Alice?*

Immediately the question arises of how *my* utterance is in fact intoned: if I too accent *Bob*, then surely my utterance would be false. One might argue that the test should assume that my report employs the nearest that the language gives

---

<sup>14</sup>See Rooth (1996) for an overview of popular semantic approaches, Schwarzschild (1999) and Geurts & van der Sandt (2004) for some recent developments in other directions.

us to a 'default' prosodic pattern. In this case, this presumably means sentence-final accenting: *In C1 and C2, Bert said that Jim only introduced Bob to ALICE.* Would this report be true? At this point, I, for one, have to admit that intuitions fail me. The report seems much more clearly *inappropriate* than either true or false. All we can say is that if there is a single proposition underlying Bert's two utterances, it isn't one that is expressible using the word string he used. This might mean that focus does contribute to minimal semantic content here, but it is also consistent with there being minimal semantic content associated with just the string of words. The fact is that we can't tell using C&L's test, because we can't utter 'just the string of words' (note that assuming a written speech report does not necessarily help, as we cannot discount the likelihood that any judge of the truth of the report mentally 'hears' one particular stress pattern upon reading it).

Examples such as these suggest 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 prejudge the nature of it for any given linguistic expression (and, consequently, inferential pragmatic processes may in principle influence all levels of meaning).

C&L really should not find this conclusion unpalatable. A good deal of C&L's apparent problem with radical contextualism (at least as represented by RT) seems to rest on terminological confusion. As I have hinted already, a number of the objections they produce against radical contextualism are in fact only problematic to a moderate contextualist position. I believe that a good deal of confusion has 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 undoubtedly 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 they share. Thus, relevance theorists tend to push the notion that contextual influence is pervasive, in contrast to the limited, grammatically triggered contextual influence found 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. But, 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. As such, these aspects of the rhetoric of RT do not contradict C&L's semantic minimalism, even though talk of con-

textual influence on propositional meanings may give that impression<sup>15</sup>. C&L's notion of proposition is knowingly removed from that of moderate contextualists (for example, 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 the latter 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>. Because we are dealing with distinct notions, it does not follow that believers in Propositionality<sub>CL</sub> have any substantial argument with RT (or any given form of radical contextualism).

RT not only does not, but could not argue for contextual influence in the creation of C&L's minimal semantic content, since the latter is, in a crucial sense, equivalent to the logical form (=encoded meaning) of an utterance in RT. Both C&L's minimal semantic content and RT's logical forms are defined as whatever meaning is common to a linguistic expression across all contexts, so there is no way in which RT's contextualism could extend this far. Therefore, as mentioned above, the only bone of contention in principle between C&L and RT is C&L's claim that minimal semantic content is propositional, in C&L's very particular sense. It would be a much better use of C&L's energies and those of relevance theorists to investigate the significance of this claim—whether it has really important consequences at this level, or whether it amounts to little more than a terminological squabble—rather than arguing at cross-purposes about whether content can be shared or whether a given lexical item is context-sensitive.

## 5.1 RT principles and moderate contextualist practice

While I argue that C&L's arguments against RT ultimately fail to undermine the framework in principle, I do believe that their broader arguments have important connotations for the practice of relevance-theoretic analysis. As I have suggested above, there is a sense in which C&L's criticisms do sometimes hit the mark. When they do, this is because of relevance theorists adopting (explicitly or otherwise) aspects of moderate contextualist thinking, specifically the idea that the logical form of a sentence—hence ultimately the lexicon or grammar—should

---

<sup>15</sup>The above-cited eighth principle of Speech Act Pluralism (C&L 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 (2006, 14\*\*\*) notes.

contain something that accounts for any given piece of contextual influence on the truth conditions intuitively conveyed by an utterance. This idea (which is, in essence, C&L's Mistaken Assumption) is (*contra* C&L 2005, 54) not intrinsic to RT. Indeed, to my mind it is anathema to the spirit of RT, far from being required by it.

At this point it is important to get away from making statements about moderate and radical contextualists, as though theoretical persuasions were like nationalities, and instead speak of moderate and radical contextualist *arguments*. The appropriateness of such labels ultimately depends upon the nature of ones theoretical claims and analytical practice, not on ones declared affiliations to a given school of thought.

C&L's only reason for lumping together moderate and radical contextualists as believers in the Mistaken Assumption (apart from a single quotation from the radical contextualist—but not relevance theorists—Charles Travis) is their common use of what C&L call Context Shifting Arguments. This refers to the strategy of showing how what is intuitively asserted by some sentence varies from context to context and on the basis of this claiming the semantics of 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? (Cappelen & Lepore 2005, 54)

There is, of course, one very practical 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 to pervasive. In other words,

the Mistaken Assumption is a premise used in a *reductio* of moderate contextualism by both C&L and radical contextualists.

C&L are therefore wrong to deduce the existence of the Mistaken Assumption from the very use of CSAs, because they fail to recognise the different uses to which CSAs are put by moderate and radical contextualists. Another way to understand their mistake 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 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 betrays the Mistaken Assumption is by definition a moderate contextualist argument. A truly radical contextualist position, such as that of RT, does not use ‘Context Shifting Arguments’ to justify positing some individual piece of context-dependent encoding and thereby to patch up a neatly compositional view of truth-conditional meaning, but rather to demonstrate the general untenability of such a view.

Where RT is concerned, this much is quite clear in the fundamentals of the framework. The starting point for Sperber & Wilson’s (1986) original exegesis of RT is the involvement of the general inferential abilities of humans in communication and the fact that there is no *a priori* way to tell just how much of what is communicated is encoded and how much inferred. Carston’s (1986) major expo-

sition of RT sets out from the rejection of Gricean notions of ‘what is said’ and of other attempts to establish intuited truth-conditional meaning as the point at which the contribution of the grammar stops and inference begins. Accordingly, when these authors use ‘Context Shifting Arguments’, they do so just in order to demonstrate how perceived truth-conditional meaning bears no particular relationship to encoded meaning.

For sure, there is nothing in all this that rules out the possibility that some given linguistic expression encodes a trigger for some context-dependent inference, as one might find in moderate contextualist analyses, but nor is this required by RT to account for context-related variation at any given level of meaning. All that the fundamentals of RT require of a relevance-theoretic analysis is that there should be a chain of reasoning from the encoded meaning (whatever form this may take) to the perceived meaning, which results from the application of the Principles of Relevance in a given context. It is a crucial tenet of RT in this respect that the same mechanisms are involved in the inferential derivation of perceived basic truth conditions as in the derivation of implicatures—a position that C&L (2005, 204) insist upon themselves. Perceived truth-conditional meaning simply does not have a special status as far as the mechanisms of RT are concerned and as such RT *cannot* involve the Mistaken Assumption.

The RT notion of *explicature* may appear to contradict my claims here. In this word, RT does have a specialised term for what we could call basic truth-conditional meaning. It is understandable that this might lead to the impression that this is a theoretically significant level of meaning in RT. But the term is essentially just descriptive: it refers to one kind of (intuitively) identifiable *outcome* of general relevance-theoretic reasoning; there remains only one mechanism underlying that reasoning. Explicature is technically defined as propositional meaning that follows from inferential development of the logical form of an utterance (in context), whereas implicature (as in more conventional approaches) involves the creation of propositional meanings that may bear little relation to the form of the utterance.

As Sperber & Wilson make clear (see in particular Wilson & Sperber 1993), there is not necessarily a single explicature associated with each utterance. This in itself makes it clear that explicature is not the kind of intuitive ‘what is said’ meaning that could be associated with the Mistaken Assumption. For example, if Rob utters *She’s eaten* in a certain context, one explicature (i.e. one development of the logical form) might be rendered as ‘Cassie has eaten lunch on 1 May 2006’, but this may lead to ‘higher-level explicatures’ such as ‘Rob believes that Cassie has eaten lunch on 1 May 2006’ and well as likely implicatures such as ‘the

addressee does not need to cater for Cassie'. All of this is the result of one process of relevance-driven inference of Rob's communicative intentions. The terms *explicature* and *implicature* should be taken as expository devices that facilitate discussion of a chain of reasoning, and of the particular contribution of linguistic material to this<sup>16</sup>. Talk of explicatures therefore in no way implies adherence to the Mistaken Assumption, which remains definitive of moderate contextualist arguments<sup>17</sup>.

It is important to note that explicature may not be seen by all relevance theorists in this way. Some may wish to afford more genuine theoretical status to the idea of explicature (a tendency that shows through in Carston's treatment of numerals, discussed immediately below). In this case, my comments above can be taken as arguments as to how RT should be interpreted, if it is to be a coherent approach to the derivation of meaning (and one that is distinct from the dominant moderate contextualist paradigm). In this respect, some relevance theorists may have something to learn from C&L: even if they fail to undermine RT in principle, they do show us how RT should not be thought of. If their arguments stick, RT is being interpreted in the wrong way. I do not believe that this is revisionism on my part, but if it is (inadvertently) so, I am happy for my arguments to be taken on that level: as a plea for a particular interpretation of RT.

Thus, while I maintain that RT is in principle the wrong target for C&L's arguments regarding Context Shifting Arguments and the Mistaken Assumption, there is certainly a sense in which many avowed relevance theorists need to bear C&L's arguments in mind. As I hope is by now clear, my own belief is that RT is a coherent and useful approach just as long it maintains the C&L-like minimalism that (I believe) is to be found in the original formulation of the theory. This is minimalism with respect to the amount of propositional meaning that should be accounted for in the lexicon and grammar, which inevitably means a heavy explanatory burden on the inferential mechanisms that RT proposes. This latter fact

---

<sup>16</sup>Consequently, any apparent inconsistency in defining explicature does no serious damage to the principle claims of RT, contrary to what Levinson (2000), Atlas (2006) seem to believe. RT stands or falls on the adequacy of the Principles of Relevance as a basis for drawing certain inferences, not on what we may choose to label certain steps along the way.

<sup>17</sup>If the notion of explicature is suspect in any way, this is because of its relationship to the idea of logical form. As I have argued elsewhere (Wedgwood 2005, 55–58), RT, as a theory of the *process* of interpretation, ought to be capable of investigating inferences that may be made in the course of parsing. As such, relevance-theoretic reasoning should not begin with a logical form, at least if this is interpreted as meaning a representation of the encoded meaning of a whole utterance. If the notion of a single, static logical form is abandoned in this way, then the notion of explicature of course must either be dropped or re-defined.

may be uncomfortable for many analysts, but it should be borne in mind that these mechanisms are independently necessary.

## 5.2 Insufficiently radical RT: moderation as a cardinal sin

One example of an RT analysis that seems to involve implicitly moderate contextualist assumptions is Carston's (1985, 1998) account of 'scalar' inferences arising from quantificational items. Carston argues against the (still widely assumed) neo-Gricean analysis of cardinal numerals Horn (1972, originating with), whereby they are taken to encode 'at least  $n$ ' readings and take on 'exactly  $n$ ' reading only by implicature. Evidence against this comes from a variety of sources. For example, Carston cites Sadock's (1984) argument that simple mathematical statements would have ludicrous literal meanings on the neo-Gricean analysis (*The square root of 9 is 3* does not mean 'the square root of at least nine is at least three'). Numerals also easily allow so-called 'scale reversals' (i.e. implicit 'up to  $n$ ' / 'at most  $n$ ' readings), as in *She can eat 2000 calories a day and not put on weight*. A significant part of Carston's argumentation concerns ways of demonstrating that inferences over numerals affect intuitive truth conditions and are therefore not mere implicatures but matters of explicature. As far as it goes, this is legitimate enough as a response to the neo-Gricean (hence moderate contextualist) line; however, Carston slips into moderate contextualist reasoning when she goes on to contrast cardinal numerals with other expressions that trigger what would traditionally be termed 'quantity' inference (after Grice's maxim).

Carston approvingly cites Horn (1992) (in turn a response to Carston 1985), who argues that data like (6) and (7) constitute evidence for a difference in kind between the cardinals and other quantificational expressions:

- (6)       A: Do you have two children?  
          B<sub>1</sub>: No, three.  
          B<sub>2</sub>: ?Yes, (in fact) three.
- (7)       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.

The idea here is to show that numerals and other quantifiers relate to truth conditions in different ways. While we know that both *two* and *many* often get upper-bounded readings ('exactly two' and 'many but not all'), the judgements

attached to these data suggest that this reading of *two* constitutes part of the truth-conditional meaning that the speaker is committed to in (6), while the upper bound on *all* is easily cancelled in the manner of a classic Gricean implicature in (7)<sup>18</sup>.

Similarly, Carston criticises data from Scharten (1997, 66–67) which purport to show that a single principle of ‘exhaustivity’ should account for quantity inferences with numerals and non-numerals:

- (8) Q: How many pupils are there in your class?  
A: 31.
- (9) Q: What is your profession?  
A: I am an architect.

Carston notes that Scharten herself appears to assign (9) a different status with regard to truth-conditionality:

She says of example [(8)] “If it turns out that there are 33 pupils in the speaker’s class, then he will not have spoken truthfully.” But in her discussion of [(9)] [...] she is far less categorical (rightly so, in my view): “Here again, if it turns out that the answerer is also a practising doctor, then she may not have spoken the whole truth [...] The answer [...] may well be called true but it is incomplete ...”. The implication seems to be that the inference (if there is one) that A is just an architect [...] does not affect the truth-conditions of the utterance; that is, it is an implicature. (Carston 1998\*\*\*)

Carston proceeds to argue for a position whereby quantity inferences over numerals *inherently* contribute to explicature, while non-numerals trigger mere implicature. Accordingly, she proposes a lexical semantics for numerals that stipulatively triggers inferences to an ‘at most *n*’, ‘at least *n*’ or ‘exactly *n*’ reading. This is to be achieved by adding a variable to the lexical semantics, to give a representation such as [X[THREE]]. 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\*\*\*)

---

<sup>18</sup>I put aside the significant question of whether *yes/no* responses should be taken as markers of truth-conditionality by any definition (see Geurts 1998).

The structure of Carston's analysis is thus a classic moderate contextualist one: first one shows that some linguistic item triggers inferences in context that affect intuitive truth conditions and then one adds something to the encoded semantics of this item so that this ability to affect truth conditions is traceable to the grammar. In adding a variable to the semantics of numerals just to trigger the observed inferences, Carston openly draws on moderate contextualist proposals for quantifier domain restriction (to be subsequently defended at length by the 'hidden indexicalists' and critics of RT, Stanley & Szabó 2000). The question is why Carston, as a relevance theorist, would adopt such a strategy.

There should be no assumption in an RT analysis that the ability (or even a regular tendency) to affect truth-conditional meaning implies some special status in semantics. RT deals in general principles of inference in context and emphasises how these are always at work in the interpretation of ostensive communication. Such inferences therefore rarely if ever require direct and explicit triggers—the addressee of an utterance doesn't need telling to seek the relevance of a numeral or of any other linguistic item.

Furthermore, when two instances of inferential reasoning appear to share a basic logic structure it should be the normal thing in RT to assume that essentially the same process is at work, irrespective of whether the outcome of either happens to be describable in terms of an effect on truth conditions. The basic notion of 'quantity' inference is rather obviously an example of such a common inferential structure: thanks to very general principles, an addressee can expect a speaker to mention everything that is relevant in a certain way<sup>19</sup>, so that anything that is manifestly potentially relevant and unmentioned can be taken to be (held by the speaker to be) not the case. It seems clear that this kind of inference is likely to have different effects—in terms of how strongly it is taken to commit the speaker to certain claims—when applied in conjunction with different kinds of assertion (e.g. involving different classes of lexical item). From the analyst's point of view, it therefore seems much more important to characterise the nature of the gen-

---

<sup>19</sup>It is easy enough to see this without getting into the technicalities of any pragmatic framework. By the very act of uttering anything, a speaker signals a willingness, and generally desire, to convey certain information—information that is relevant in a certain way. If the speaker is thus signalling willingness to share *that which is relevant in a certain way*, then (*ceteris paribus*) there is only one inferable reason for leaving out some information that is potentially relevant in that way: the speaker does not hold it to be so. Of course, speakers may withhold information deceitfully, in which case they rely on this same inference being drawn, or openly, in which case they effectively signal that the withheld information is in fact irrelevant in *other* ways. Note that no overarching Gricean drive towards co-operation need be involved in explaining 'quantity' inferences.

eral inferential process than to worry about whether or not it happens to produce what we could call a truth-conditional effect in certain circumstances. RT rightly recognises that general inferential processes can be of broad explanatory power and rejects the Gricean notion that they cannot be involved in deriving perceived truth conditions. As such, the relevance theorist simply has no motivation to account for explicature/implicature contrasts in lexically or grammatically encoded semantics.

Should the analyst nevertheless be keen to explain the observed tendencies of certain lexical items to contribute regularly to explicature rather than producing implicature, the logical question to ask is *how* the encoded semantics of different classes of item interact with the Principles of Relevance, not whether they do obligatorily or otherwise. There are obviously several important differences between numerals and other expressions; any one of these (or a combination of these) might affect the degree to which a speaker is typically perceived to be committed to an inferred ‘upper bounded’ reading.

It is not my aim here to provide a full analysis of the differences between numerals and non-numerals, but consider the following fact: numerals are inherently scalar in a particularly strict way. The very definition of a given cardinal value depends upon its place in the scale of integers. Quantifiers like *many*, on the other hand, are not nearly so inherently scalar (even if one *can* identify an entailment from *all* to *many*, one might define *many* without reference to this entailment), nor are they so precisely defined. Non-quantificational items like *architect* are plainly still less inherently related to other items—a particular context may include an ordered set of alternatives against which such an item may be interpreted exhaustively (Hirschberg 1991), but on the whole such a term commits a speaker to the exclusion of other possibilities only in a rather general way, and consequently rather weakly. It is easy to see that if item A is *inherently* defined in contrast to item B (perhaps among others), use of A will imply the rejection of B rather regularly and rather strongly. Also, the more precisely defined some item is, the more a deliberate choice to employ that item implies the rejection of alternatives to it (including higher values on a scale); so *six* is more likely to commit the speaker to ‘not seven’ than *many* is to commit the speaker to ‘not all’. Therefore, even on a rather conventional view of ‘scalar’ inferences (i.e. as involving an inferred upper bound), there are uncontroversial aspects of the differences between numerals and non-numerals that should lead us to expect the former to commit the speaker more readily to an upper bounded reading—and thereby produce non-cancellable, hence ‘truth-conditional’ meaning of this kind—without assuming anything other than a very general kind of ‘quantity’ inference.

The rights and wrongs of particular analyses of scalar inference need not concern us here, however. In fact, there are good reasons—for what it's worth—to reject both the generalisation that non-numerals only create implicatures (Wedgwood 2005, Chapter 5) and the idea that upper bounds are necessarily inferred (see Koenig 1993, Kadmon 2001, Breheny 2005, Wedgwood 2005). The point of the previous paragraph is merely to show that there is no *need* to account for the perceived truth-conditionality of upper bounds on numerals by brute force (if one feels the need to account for it at all). The bigger issue on a theoretical and methodological level is that to even think of doing so runs counter to the spirit of RT, which is distinguished from its Gricean predecessors by the very fact that it does not grant special status to the derivation of perceived truth-conditional meaning (i.e. it does not make the Mistaken Assumption). 'Quantity' inferences are predicted to occur by RT as part of the overall workings of the Principles of Relevance (Sperber & Wilson 1995, 272–278), and as such should be seen to occur without special triggering by certain classes of lexical item. An inference that is made obligatory by a variable in logical form is simply not the same kind of relevance-driven inference. It is rather some arbitrary piece of compositional semantic encoding, and furthermore one that exists just to support a theoretical distinction that RT in principle does not care about<sup>20</sup>.

That Carston of all people should propose such an analysis is surprising, since she elsewhere argues forcefully for the idea that elements of perceived propositional meaning can be inferred rather freely, without a lexically or grammatically encoded trigger. For example, she rejects the idea of Stanley (2000) and Stanley & Szabó (2000) that the existence of a bound variable in some communicated propositional meaning constitutes evidence for a variable in logical form (Carston 2002, 200). In so doing, she is defending a truly radical contextualist (and therefore more coherent) reading of RT; one that is free of the Mistaken Assumption and thus immune to C&L's criticisms. Whether this reflects a conscious change in Carston's thinking since her (1998) article on numerals, or a degree of theoretical inconsistency in her work, is unclear.

A still more general point lurks behind all this. 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

---

<sup>20</sup>This is not to say that what is inferred in one place cannot be encoded in another. This is common enough cross-linguistically and language-internally, synchronically and diachronically. But this concerns the difference between inferred material and material that is directly encoded in the lexicon or grammar (a distinction RT that does very much recognise), not a difference between freely inferred material and material inferred on the basis of an encoded trigger.

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, this is essentially the same break that C&L make. Relevance theorists should be no more afraid of taking a radical stance than they are<sup>21</sup>.

Above all, believers in these approaches should accept 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 makes it impossible simply to assume conventional semantic analysis, so much of which is based on the compositional derivation of perceived truth conditions. Once we reject this, as RT does in principle, the nature of encoded meaning becomes an entirely open question. The same applies to C&L's approach, assuming that ultimately we seek slightly more enlightenment on this score than what we can glean from circular T-sentences like *'Steel isn't strong enough' is true iff steel isn't strong enough*. 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 relevance theorists and C&L 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). RT at least gives the basis for future analysis by proposing a set of principles underlying the crucial inferential component in linguistic interpretation; C&L fail to do this and as such fail to provide a basis for following through the consequences of their own arguments in actual linguistic analysis.

---

<sup>21</sup>C&L make out that they are not radical but positively conservative in their 'semantic minimalism' (2006b\*\*\*), but elsewhere make clear that they are swimming against the tide: "Hardly any contemporary philosopher rejects [moderate contextualism]" (2005, 88).

## 6 C&L's misapprehensions about RT

As noted above, C&L (2005, 181) know that some people see a likeness between them and relevance theorists and they are at pains to point out that they are quite different. Their list of differences does at least home in on what I have identified as the one noteworthy difference between C&L's ideas and RT, the claimed propositionality of encoded meaning (or minimal semantic content). All but the final point in the list relate directly to this; however, the implications that C&L draw from it are simply not justified by the nature of their position more generally.

Their list is nothing if not pithy; I shall attempt similar concision in responding to each point on it. While doing so involves some repetition of arguments made above, this thereby provides a useful summary of these arguments and helps to identify just where C&L's objections to RT break down.

- (a) "We think, and the Relevance Theorists deny, that there is a minimal semantic content or proposition that is semantically expressed by (almost) every utterance of a well-formed English sentence [...] a full-blooded proposition with truth conditions and a truth value. [...] This is a substantial disagreement about the metaphysics of content."

This point has largely been dealt with above. The claim that encoded meaning is propositional is undoubtedly a real difference between C&L and RT. On the other hand, C&L's notion of proposition here is something quite different to the kind of 'truth-conditional meaning' that relevance theorists are used to arguing about (the moderate contextualists' intuited truth conditions). Hence RT's denial of fully encoded propositions is hardly relevant to C&L's claims. By their own arguments, what C&L are happy to call propositions are regularly what moderate contextualists would call incomplete or underdetermined, and it is in this context that RT's claims should be assessed. *In the sense in which the term is typically used*, encoded meanings are indeed 'sub-propositional', according to both RT and C&L. The extent to which propositionality in C&L's sense really matters (i.e. whether we stand to lose anything in practical terms by simply talking of encoded meaning), and whether it is compatible with RT, are to my mind open questions—though I have indicated that C&L's linguistic tests for propositional minimal semantic content are not reliable.

- (b) "We think, and the Relevance Theorists deny, that this minimal semantic content is an essential part of all communicative interactions. The minimal

semantic content has a function in the cognitive life of communicators that no other content can serve.”

- (c) “We think, and the Relevance Theorists deny, that minimal semantic content has a psychological role that no other ‘level of content’ can fill.”

Replace the phrase *minimal semantic content* with *encoded meaning* or *logical form* and these ideas constitute an integral part of RT, not something its practitioners would ever deny. The minimal sense in which C&L’s minimal semantic content is propositional plays no role in its being ‘an essential part of all communicative interactions’. There is no way that C&L could argue that this depends on propositionality, since the detailed implications of their kind of propositionality are left quite undefined (truth conditions at this level being avowedly inaccessible to semantics, beyond an essentially circular statement like (1)). The crucial cognitive function / psychological role served by minimal semantic content is just the provision of a suitable basis for calculating intended (SA) meanings, given some independent set of principles for making such calculations. This is precisely the function of encoded meaning in RT. So there is no substantive difference between C&L and relevance theorists here.

- (d) “We think, and the Relevance Theorists deny, that theories that do not recognize minimal semantic content are empirically inadequate and internally inconsistent.”

This claim relates to all of the arguments presented above. RT might be less empirically inadequate and internally inconsistent than C&L’s position if minimal semantic content were responsible for what is said/asserted/claimed/etc. As we have seen, it is the ability to explain how addressees recover this kind of meaning that determines the empirical adequacy of a given theory of linguistic communication (in terms of explaining assessments of assertions, collective deliberation, and so on)—and a theory is certainly internally inconsistent if it renders recognition of this kind of meaning impossible. But C&L have misunderstood RT and possibly the logic of their own position if they think that the propositionality or otherwise of encoded meaning has a bearing on this. By their own arguments, ‘what is said/asserted/claimed/etc.’ is a description not of minimal semantic content, but of SA content. Again by their own arguments, SA content is not directly communicated but is conveyed via mechanisms that also convey implicatures. The claim that encoded meaning is propositional does not in itself determine the nature of such a mechanism and as such has no direct relationship to the empirical adequacy

or internal consistency of any theory. Given the indirectness of the recovery of SA content under both approaches, any argument from content-sharing that would automatically render RT empirically inadequate or internally inconsistent would do the same for C&L's semantic minimalism.

- (e) "Finally, [...] we emphatically reject theories that are 'Original Utterance Centrist'. All Relevance Theorists are Original Utterance Centrists."

This issue is somewhat different to (a)–(d), and less central to the arguments made in the present article. C&L's comments on Original Utterance Centrist (OUC) are rather elliptical (2005, 201–202), but suffice to suggest a genuine departure from anything a relevance theorist would endorse—one that strikes me as bordering on the bizarre, but which is perhaps to be understood with reference to the discussion of objective versus subjective assessments of communicated meaning earlier in this article.

OUC is described as the idea that "speech act content is fixed by facts about the speaker, his audience, and their common context". Contrary to this idea, C&L "think facts not known or available to the speaker (or his audience) can make a difference" (2005, 201). The obvious question here is 'a difference to what?'. The only sensible answer to this would seem to be 'a difference to what could be reported by a third party to have been communicated'. This is certainly not something that RT would claim to account for, since RT is an attempt to explain how *addressees* recover speakers' intended meanings (to the extent that they do)—which is also notably the basis on which C&L (2006a) attempt to criticise RT.

Consider C&L's example, as follows:

[...] suppose you uttered ([8]) several weeks ago:

([8]) The table is covered with books.

Suppose that whichever table is under discussion currently sits comfortably in your father's office (although it did not sit there when you uttered ([8])). Haven't you said with your utterance of ([8]) that the table in your father's office is covered with books? (Cappelen & Lepore 2005, 201)

What we have here is of course some relation of classic problems of intensionality (along the lines of knowing Smith's murderer is insane). Happily, theories of utterance interpretation like RT simply do not need to get involved in these issues—and, by the same token, nor do C&L, to the extent that they are proposing

an alternative to approaches like RT. Should some third party indeed report that *Dan said that the table in his father's office is covered with books*, this would be a new utterance and the process of recognising its meaning would be a different one to that of interpreting the original utterance. It would proceed according to the same mechanisms, but in a different context, which means against the background of a different set of assumptions. The ability to associate 'new' content with a previous utterance in the act of reporting it is thus entirely unsurprising from the point of view of RT or any comparable theory, and is not something that such a theory need have anything special to say about. 'Original Utterance Centric' seems an oddly redundant way to label such a theory, as it is precisely a theory of how any given utterance is understood by its addressee. None of the concerns that C&L express elsewhere about frameworks like RT depend on the ability of these frameworks to do anything else. As such, C&L's concerns about OUC seem to represent a sudden and unmotivated shift of focus, from SA content *qua* what is communicated to the addressee of an utterance to some notion of SA content *qua* what could be reported by a third party to have been communicated. I take the issue to be fundamentally irrelevant to current concerns, except that it provides another illustration of the importance of distinguishing the addressee's subjective perspective from an objective third party perspective: these play crucially different roles in the theory of language interpretation and therefore any failure to distinguish them risks the creation of significant confusion.

## 7 *happy to infer content?*

A brief look at one final relevance-theoretic idea mentioned by C&L (2006a, 19\*\*\*) will provide a kind of summary to this article, as well as providing a telling concrete illustration of the central issues. This is the idea of *ad hoc concept construction*, a notion promoted by Carston (see in particular 2002, Chapter 5) to account for the sometimes radical variation in meaning that can be conveyed by a single lexical item in different contexts (including various kinds of polysemy, 'loose use' and metaphorical usage). According to C&L, this is a telling example of just how radical and problematic the consequences of RT are: just how little content would be shared if RT were true. Yet (in line with my arguments above) Carston seems in fact to admit of no more variation in meaning than C&L's semantic minimalism commits them to. In effect, Carston's *ad hoc* concepts are directly equivalent to lexical semantics as viewed from C&L's metaphysics-free perspective. Again, the argument from 'non-shared content' is a red herring, and

again one can identify disagreement between C&L and RT as arising from simply talking at cross-purposes as they attack moderate contextualist assumptions from different perspectives.

C&L quote one of Carston's discussions of the adjective *happy*:

In one context, an utterance of *I'm happy* could communicate that the speaker feels herself to be in a steady state of low-key well-being, in another that she is experiencing a moment of intense joy, in yet another that she is satisfied with the outcome of some negotiation, and so on. The general concept HAPPY encoded by the lexical item *happy* gives access to an indefinite number of more specific concepts, recoverable in particular contexts by relevance-driven inference. (Carston 2004, 11\*\*\*)

They go on to criticise this view as follows:

There is, on Carston's view, an indefinite number of concepts that could be communicated by an utterance of a sentence containing *happy*. She doesn't tell us how many 'an indefinite number' is, but it's probably a lot. Which one an interpreter latches on to, depends on what is relevant to that interpreter and that depends on which cognitive effects the utterance has on her. That, again, depends on what beliefs and other cognitive states she's in at the time of interpreting the utterance.

In conclusion, Were RT true, it would turn out to be a minor miracle if two interpreters were to end up with the same interpretation of an utterance of 'I'm happy'. (Cappelen & Lepore 2006a, 19\*\*\*)

Putting aside C&L's misrepresentation of RT—the cognitive effects derived associated with an utterance by an addressee depend primarily not on arbitrary 'beliefs and other cognitive states she's in at the time', but rather on the evidence she has for the beliefs and assumptions of the speaker at the time—consider how C&L's own treatment of *I'm happy* would go:

C&L would presumably state the minimal semantic content of the sentence as follows.

- (10) 'I'm happy' expresses the proposition *that [the referent of the indexical I in the original context of utterance] is happy* and is true iff [the referent of the indexical *I* in the original context of utterance] is happy.

The precise *conditions* under which this is true would be declared to be not a matter for semantics. Meanwhile, what the speaker is taken to have said/claimed/asserted (etc.) by a given utterance of *I'm happy* would be a matter of SA content, which is derived by the same mechanisms as implicatures.

Thus, the particular contribution of the word *happy* under this view amounts to the following. It makes some contribution to the composition of a meaning that is not identical to what is said/claimed/asserted (etc.) by an utterance of *I'm happy*, whose precise truth conditions are unknowable to semantics alone. At the level of SA content, (thanks to the involvement of the processes that also produce implicatures) its use may result in a wide variety of detailed interpretations, in different contexts.

The striking thing about the previous paragraph is that it would serve as a decent description of Carston's position without changing a word (though one would normally swap the odd bit of terminology). Carston's *ad hoc* concept corresponds directly to C&L's minimal semantic content for *happy*: it is some encoded meaning, to which the semanticist has no direct access, that is the common basis for a variety of assertions/claims/etc. once it is set in context and subject to some general mechanism for divining intended meanings. Note once more that because C&L, like RT, take this mechanism to be common to the derivation of 'what is said' and 'what is implicated', they must assume it to be based in assessments of identifiably shared (in RT terms, mutually manifest) beliefs (assumptions) and to be fallible; they surely do not want to claim that every implicature is infallibly recovered and that communication in this sense never breaks down<sup>22</sup>.

Therefore, by C&L's (2006a) arguments, it should be "a minor miracle if two interpreters were to end up with the same interpretation of an utterance of 'I'm happy'" if interpretation proceeds according to their own model. The extent to which interlocutors 'share content' and how they do so are questions that relevance theorists and C&L both have to face, in the same way and for the same reasons. RT, unlike C&L, does propose an answer. Whether it is a good answer or not is of course open to debate, but not on the grounds of the impossibility of 'sharing content'.

Finally, consider the point of discussing word meanings in terms of *ad hoc* concepts. What point is made by emphasising the degree of semantic variation associated with a single word? It seems quite clear to me that in pushing this notion, Carston is taking on the idea that perceived truth-conditional meaning—i.e.

---

<sup>22</sup>Elsewhere, C&L do come close to explicitly acknowledging the fallibility of communication (2005, 206)

what is said/asserted/etc.—is largely composed in the grammar. By emphasising the context-sensitivity of simple predicates like *happy* (and many other words), Carston shows the pervasiveness of context-dependent inference in the creation of communicated meanings. In other words, the idea of *ad hoc* concepts emphasises two of C&L's central beliefs: (i) that what is said/asserted/etc. depends on the same mechanisms that produce implicatures and (ii) that contextual influence on this level of meaning cannot be restricted to certain classes of item (i.e. moderate contextualism is unsustainable).

So how do C&L come to the conclusion that *ad hoc* concept construction represents something quite opposed to their position, with significantly different implications? I can only assume that Carston's particular way of presenting the idea in question leads them to this. Here one may have some sympathy with C&L, given that they approach Carston's writings from a philosopher's perspective, not that of a pragmaticist linguist. Carston's underlying point could be made in more conventional terms, by once again simply appealing to a recognised tool of formal semantics, the denotation assignment function. If semanticists were properly mindful of the assignment function that each and every expression of a semantic representation language carries (explicitly or implicitly), moderate contextualism would never come about. For the ubiquity of assignment functions implicitly relativises everything to context. To say that *happy* makes a different contribution to truth conditions in different contexts is in effect merely to state the semantic truism that contribution of *happy* depends upon the particular function that assigns a denotation to  $\llbracket happy \rrbracket$ . I am not suggesting that Carston had this in mind when proposing *ad hoc* concept construction, but it is logically the case.

One might forgive C&L for failing to see that this is in effect Carston's point, since she expresses it in terms that are so distinct from conventional semantic theory. On the other hand, if the long-standing existence of assignment functions in semantic theory doesn't suffice to ensure that semanticists avoid moderate contextualism, perhaps it is worth putting the argument differently. Furthermore, Carston's way of framing the issue does something tremendously important, from the point of view of the relevance theorist or speech act pluralist: it draws attention to the very *active* role of context-dependent inferential processes in determining what is perceived to be asserted.

In conclusion, C&L share much more with RT than they, or most relevance theorists, seem to realise. Where RT has encoded meanings, C&L have minimal semantic content. Where RT has multiple explicatures and implicatures, C&L have speech act pluralism. Both take as their starting point a rejection of the idea that grammatically encoded meaning should systematically account for 'what is

said'. There remain just three identifiable differences between them, only one of which is clearly a point of principle. 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, C&L speak of rejecting Original Utterance Centrism, an issue that seems entirely dependent on their (in my view inappropriate) concentration on objective, third party assessments of utterance meaning. Third, the one clear difference of principle, C&L insist that minimal semantic content is propositional, while RT makes no *a priori* claims for encoded meaning.

Whether or not C&L could actually be called disguised relevance theorists depends on the weight one gives to this latter issue. The extent to which the propositionality of minimal semantic content constitutes a difference of real consequence is a matter that may be left to future discussion. What certainly does not distinguish the two approaches, let alone undermine RT, is the argument that RT makes it impossible to share content. Instead of attempting to distance themselves from each other by such misguided arguments, these two schools of thought have a good deal to gain from joining forces. C&L need a theory of how pragmatic inference, and hence a large chunk of linguistic interpretation, actually works, and this is just what RT is. Meanwhile, relevance theorists have something to learn from C&L's uncompromisingly minimalist stance.

## References

- Asher, N. & Lascarides, A. (2003), *Logics of Conversation*, Cambridge University Press, Cambridge.
- Atlas, J. D. (2006), How insensitive can you be? Meanings, propositions, context, and semantical underdeterminacy, in G. Preyer, ed., '\*\*\*', \*\*\*, \*\*\*, pp. \*\*\*–\*\*\*.
- Austin, J. L. (1962), *How to do things with words*, Clarendon Press, Oxford.
- Blakemore, D. (1992), *Understanding Utterances: an Introduction to Pragmatics*, Blackwell, Oxford.
- Breheny, 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', NCS, 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), Relevance theory and shared content, in

- N. Burton-Roberts, ed., 'Advances in Pragmatics', Palgrave Macmillan, London. Forthcoming.
- Cappelen, H. & Lepore, E. (2006*b*), 'Response', *Mind and Language* **21**(1), 50–73.
- 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.
- Fodor, J. D. (1987), *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*, M.I.T. Press, Cambridge, Mass.
- Geurts, B. (1998), Scalars, in P. Ludewig & B. Geurts, eds, 'Lexicalische Semantik aus kognitiver Sicht', Gunter Narr, Tübingen, pp. 95–118.
- Geurts, B. & van der Sandt, R. (2004), 'Interpreting focus', *Theoretical Linguistics* **30**, 1–44.
- Heim, I. & Kratzer, A. (1998), *Semantics in Generative Grammar*, Blackwell, Malden, Mass. and Oxford.
- Hirschberg, J. (1991), *A Theory of Scalar Implicature*, Garland, New York.
- 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.
- Jackendoff, R. (2002), *Foundations of Language: Brain, Meaning, Grammar, Evolution*, Oxford University Press, Oxford.
- Kadmon, N. (2001), *Formal Pragmatics: Semantics, Pragmatics, Presupposition, and Focus*, Blackwell, Malden, Mass. and Oxford.
- Kempson, R. (1988), Introduction\*\*\*, in R. M. Kempson, ed., 'Mental Representations: The Interface between Language and Reality', Cambridge University Press, Cambridge, pp. \*\*\*\_\*\*\*.
- Koenig, J.-P. (1993), Scalar predicates and negation: Punctual semantics and interval interpretations, in L. Dobrin, L. Nichols & R. Rodríguez, eds, 'CLS27: Papers from the 27th Regional Meeting of the Chicago Linguistic Society, 1991.

- Part Two: The Parasession on Negation', Chicago Linguistic Society, Chicago, pp. 140–155.
- Levinson, S. (2000), *Presumptive Meanings: The Theory of Generalized Conversational Implicature*, M.I.T. Press, Cambridge, Mass.
- MacKenzie, I. (2002), *Paradigms of Reading: Relevance Theory and Deconstruction*, Palgrave Macmillan, Basingstoke.
- Rooth, M. (1996), Focus, in S. Lappin, ed., 'The Handbook of Contemporary Semantic Theory', Blackwell, Oxford, pp. 271–297.
- Sadock, J. (1984), Whither radical pragmatics?, in D. Schiffrin, ed., 'Meaning, Form and Use in Context: Linguistic Applications', Georgetown University Press, Washington, pp. 139–149.
- Scharten, R. (1997), Exhaustive interpretation: A discourse-semantic account. Ph.D. thesis, Katholieke Universiteit Nijmegen.
- Schwarzschild, R. (1999), 'Givenness, avoidf and other constraints on the placement of accent', *Natural Language Semantics* 7, 141–177.
- Sperber, D. & Wilson, D. (1986), *Relevance: Communication and Cognition*, Blackwell, Oxford. (Second Edition: 1995).
- Sperber, D. & Wilson, D. (1995), Postface, in 'Relevance: Communication and Cognition', second edn, Blackwell, Oxford.
- Stanley, J. (2000), 'Context and logical form', *Linguistics and Philosophy* 23, 391–434.
- Stanley, J. & Szabó, Z. G. (2000), 'On quantifier domain restriction', *Mind and Language* 15(2–3), 219–261.
- Wedgwood, D. (2005), *Shifting the Focus: From Static Structures to the Dynamics of Interpretation*, Elsevier, Oxford.
- Wilson, D. & Sperber, D. (1993), 'Linguistic form and relevance', *Lingua* 90(1/2), 1–25.
- Wilson, D. & Sperber, D. (2004), Relevance Theory, in L. Horn & G. Ward, eds, 'The Handbook of Pragmatics', Blackwell, Oxford, pp. 607–632.