

7

relevance theory and shared content

herman cappelen and ernie lepore

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. Relevance theory (RT), however, implies that content sharing is impossible; or at least, we will argue as much in what follows.

This chapter is divided into three sections. In section 1, we amplify on what we mean by 'shared content' and its roles in how we think about language and communication; we discuss various strategies RT might invoke to account for shared content and why all these strategies fail. Section 2 is exegetical; there we show why RT *must* deny the possibility of shared content. The denial is a direct consequence of some of the most central tenets of RT. It is, however, a consequence downplayed by RT proponents. Our goal in section 2 is to show how central the denial of shared content is to RT. In section 3 we outline how we think a pragmatic theory should account for shared content.

question: good enough?

According to RT, interpreters follow 'the least effort strategy' (LES):

(LES) Check interpretive hypotheses in order of their accessibility, that is, follow a path of least effort, until an interpretation which satisfies the expectation of relevance is found; then stop. (Carston 2001, 6)

In section 2 we elaborate on how exactly to interpret (LES) and its RT defence. For now, we assume (LES) and ask: How does pursuing (LES)

help the interpreter work out which proposition a speaker intends to communicate (i.e. which proposition she endorses, believes, etc.)? Here's Carston's reply:

It [i.e. the procedure described in (LES)] provides a reliable, though by no means foolproof, means of inferring a speaker's meaning. As a patently non-demonstrative inference process, it sometimes fails and doesn't come up with the intended meaning. And when it is successful what is achieved is *seldom* a perfect replication in the hearer's mind of the very assumptions the speaker intended to communicate. 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 2001, 825; emphasis added)

Carston's view then is this: A speaker utters a sentence *S* intending to communicate the proposition that *q*; the interpreter 'typically' ends up with a range of propositions *p*₁... *p*_{*n*}, none of which is identical to *q*. But that's no problem, says Carston, because as long as *p*₁...*p*_{*n*} are similar to *q*, then that *is good enough*. In this respect, Carson agrees with Bezuidenhout:

Since utterance interpretation is always in the first place colored by one's own cognitive perspective, we think we should reject the idea that there is an intermediate stage in communication which involves the recovery of some content shared by speaker and listener and which is attributed by the listener to the utterance. In communication...[w]e *need recognize only speaker-relative utterance content and listener-relative utterance content and a relation of similarity holding between these two contents* ...This does not mean that we have to deny that literal interpretation requires the preservation of something. But this something need simply be a relevant degree of similarity between the thought expressed by the speaker and the thought expressed by the listener. (Bezuidenhout 2002, 212–13; emphasis added)

Sperber and Wilson, in their classic work (1986), do not make this explicit, but they do say:

... 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 and Wilson 1986, 192–3; emphasis added)

All these writers are committed to a version of the following non-shared content principle (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*.

According to Carston and Bezuidenhout, the *R*-relation is similarity. In section 2, we argue that this is too optimistic; there's no non-trivial sense of 'similarity' in which the explicatures arrived at by using the Principle of Optimal Relevance will be similar to the proposition that the speaker intended to communicate. They will, if RT is correct, be developments of the same logical form *LF*, but an *LF* can be developed into *radically different* propositions, and there is, for example, no guarantee that the different developments will even have the same truth-value, much less the same truth conditions. So, only in a trivial sense are we, according to RT, guaranteed similarity between, on the one hand, the proposition the speaker intends to communicate and, on the other, the interpretation the interpreter will come up with by following the procedures RT claims we as a matter of fact do follow.

We leave the question of how to understand the *R*-relation for section 2. For now, we assume that *R*-relation is some kind of similarity relation – since that is common ground between RT proponents and us (though the arguments we use in section 1 have even greater force if the *R*-relation isn't even similarity, in any interesting sense).

We'll argue that (NSC) is not, in Carston's terminology, *good enough*. We have three main arguments:

- a. Theories that incorporate (NSC) fail to account for our reporting practices, i.e. our practice of reporting on what others' say (assert, claim, suggest, propose, etc.). We consider two main strategies RT can exploit to explain what we do in such reports, and both, we argue, fail.
- b. Theories that incorporate (NSC) fail to explain our practice of assessing what others assert (claim, suggest, propose, conjecture, etc): if we (typically) don't grasp the proposition that they intended to express

(but only one sufficiently similar), we can't evaluate the proposition they intended to express. At best, we will evaluate some proposition similar to the one they are committed to, and not the one they are actually committed to.

- c. Theories that incorporate (NSC) create self-referential problems – they are what we will call 'Communicatively Self-Defeating' – a view is communicatively self-defeating if its truth implies that it can't be communicated (and hence, not evaluated).

Before moving to criticism, we make one more introductory remark: RT isn't the only theory that endorses NSC. It is a view in common among all theorists that Cappelen and Lepore (2005) call *Radical Contextualists*. Slightly more precisely, it's a feature of every theory that adheres to some version of (RC1)–(RC3):

(RC1): No English sentence *S* ever semantically expresses a proposition. Any semantic value assigned to a sentence *S* can be no more than a propositional fragment (or radical), where the hallmark of a propositional fragment (or radical) is that it does not determine a set of truth conditions, and hence, cannot take a truth-value.

(RC2): Context sensitivity is ubiquitous in this sense: Fixing for linguistic context sensitivity will never, no matter how widespread, issue in more than a propositional fragment.

(RC3): Only an utterance can semantically express a complete proposition, have a truth condition, and so, take a truth-value.

Searle (1978, 1980), Travis (1985, 1989, 1996), Recanati (2001, 2004), and a whole range of neo-Wittgensteinians, are all Radical Contextualists in this sense. Many of the objections raised against RT below apply also to all Radical Contextualists, but our focus here is on the specific version of NSC found in RT.

1. relevance theory vs shared content

(NSC) is treated as an interesting corollary of RT – it's presented more or less as an afterthought, 'Oh, by the way, if you buy into our story, you'll have to give up that old fashioned Fregean idea that in successful communication we share thoughts.' We suspect one reason for this is that (NSC) is seen as a *philosophical thesis* with *no direct empirical import*. But that's a serious error. (NSC) has direct consequence for how we should expect speakers to behave linguistically. In particular, (NSC) requires RT

proponents to provide an explanation of what speakers do (how we as theorists should understand what they do) when they report on each other's speech, when they attribute beliefs to each other, and when they evaluate each other's speech and beliefs.

We discuss these in turn.

data 1: speech reports

Imagine someone, Naomi, uttering a sentence, S. To say what she said, we paradigmatically utter sentences like (1):

(1) Naomi said that p.

We end up having the belief expressed by such sentences because we interpret Naomi's utterance as having asserted that p. 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. That is, our utterances of sentences like (1) will attribute to Naomi the endorsement of a proposition different from the one she originally intended to communicate. A proponent of (NSC) has to tell us how to understand this kind of metalinguistic activity. She has but two options: (a) she can say that reports like (1) are literally false, or (b) she can say that (1) can be true even though Naomi did not intend to communicate the proposition that p. Both options are problematic for reasons discussed below.

To see the scope of this problem, consider sentences like (2) (with 'that' demonstrating Naomi's utterance of S):

(2) That's what Sally said too.

According to RT, what we've been doing is the following: we've reached an interpretation of Naomi that yields a proposition similar (but not identical) to the one Naomi intended to communicate; we've reached an interpretation of Sally that yields a proposition similar (but not identical) to the one that Sally intended to communicate, and in uttering (2), we're claiming that there's a single proposition such that both Sally and Naomi asserted it. Again, (NSC) proponents must tell us whether (NSC) implies that sentences like (2) (and more generally claims about speakers having made the same assertoric commitments) are false or should be interpreted in such a way that they come out true. Both options are criticized below.

data 2: from assertion to belief attributions

Often we exploit our conclusions about what people say in order to attribute beliefs to them. We're not mind readers, so we go from asserted contents to beliefs. In at least some context, the following is true: if we believe (1) to be true (and we also believe that Naomi was sincere, etc.), then we infer that Naomi believes that p. So in some contexts, the truth of (1) and (2), will enable us to infer that Naomi and Sally both believe that p.

If (NSC) were correct, this procedure would, at a minimum, be questionable. Remember, the proposition that we end up with in interpretation, is typically not the proposition that the speaker intended to communicate. It is a proposition similar to the one that she intended to communicate. We certainly do not believe all of the propositions similar to the ones we believe. (If we did, we would believe everything.) So, if (NSC) were true, then our practice of (sometimes) going from asserted content to belief attributions would be in jeopardy. Again, two strategies present themselves to the (NSC) proponent: (a) Agree that the attributions we end up with using this procedure are false, or (b) claim that belief reports – 'A believes that p' – can be true even though A doesn't endorse p, but only a proposition relevantly similar to p. Both of these options are evaluated below. At this point we don't mean to rule out either of these options. We just want to earmark that endorsing (NSC) has implications and so, its defence requires work.

data 3: assessments of assertions

There's a pretheoretic presumption, we think, in favour of the view that in speech and belief reports we share contents with the reportee. That presumption is even more salient in our evaluations of others' assertoric commitments and beliefs. Consider sentences like (3)–(6) (where the 'that' in (4)–(6) demonstrates Naomi's utterance):

- (3) I agree with Naomi.
- (4) That's true.
- (5) That's questionable/inflammatory/unacceptable/irrational/clever.
- (6) There's no evidence for that.

Again, there's tension between (NSC) and (3)–(6):

- There's tension between (3)–(4) and (NSC) because in uttering either (3) or (4) the speaker doesn't mean to express agreement just with a proposition similar to the one endorsed by Naomi's utterance;

the speaker of (3) or (4) means to express agreement with the very proposition/thought that Naomi expressed.

- There's a tension between (5)–(6) and (NSC) because in uttering (5) or (6) the speaker doesn't mean to evaluate a proposition similar to the one asserted by Naomi's utterance; the speaker of (6) or (7) means to evaluate the very proposition/thought that Naomi expressed.

1.1 revisionism and conservatism

The challenge for (NSC) proponents is to explain what the implications of their theory are for our practices of describing others' speech act contents, of going from such descriptions to belief attributions, and then, for evaluations of the attributed contents. We imagine (NSC) proponents responding to this data in one of two ways. (These strategies mirror the responses philosophers who advocate apparently counter-intuitive positions in other areas of philosophy make use of.)

1. Revisionist Strategy: RT implies that the statements we make (and the thoughts we have) about shared content (described in (1)–(6) above) are false. When we utter sentences like (1)–(6) (or think the thoughts they express), we make false claims (and have false thoughts). The Revisionist Strategy can be compared to an error theory about moral statements: Our moral language presupposes there are moral facts. It turns out that there are none. So, what we say when we make moral claims is false. Similarly: The concepts we use for talking about content presuppose that contents can be shared across contexts, but, according to the Revisionists, contents aren't shared across contexts. So, we end up making false statements and having false thoughts about content. Normal speakers not initiated into RT think that they can duplicate thoughts and that successful communication requires matching of propositions, but they are entirely wrong. RT reveals this pretheoretic prejudice for what it is.
2. Conservative Strategy: The Conservative Strategy says that when correctly interpreted, Data 1–Data 3 are *not* incompatible with (NSC). It's a confusion, the Conservative says, to think our common sense beliefs about shared content are incompatible with (NSC). When these pretheoretic beliefs (and the sentences used to express them, like (1)–(6)) are properly interpreted, they are, on the whole, true. The particular version of this we consider below says that utterances of the form 'A said that p' are true just in case A said something similar to p. That, the Conservative points out, is compatible with (NSC).

The discussion of these strategies has three parts: First, we present some objections to the Revisionist Strategy; then some objections to the Conservative Strategy. These first objections focus on how the two strategies deal with speech and belief reports. Then finally we discuss the ways in which both positions fail to provide an adequate account of our practice of evaluating others' sayings and beliefs.

1.2 objections to the revisionist strategy

Revisionists, as we have described them, imagine their revisionism limited to certain meta-linguistic sentences and thoughts, those instantiated by (1)–(6). In their judgement, rejection of our belief in (1)–(6) is not too high a price to pay for a psychologically adequate theory of interpretation. The Revisionist thinks of herself as eliminating a non-theoretical dogma about content – a Fregean myth the denial of which has very little cost.

That impression is, however, quite mistaken. Our practice of content sharing is inextricably intertwined with other practices that figure centrally in our non-linguistic lives. If our beliefs about shared content are false, as the Revisionist claims, then our basic foundational beliefs about, and our understanding of, these non-linguistic practices are also in jeopardy. Here's what we have in mind:

- Coordinated action: Often, people in different contexts are asked to do the same thing, e.g. pay taxes. They receive the same instructions, are bound by the same rules, the same laws and conventions. For such instructions to function, we must assume a wide range of utterances express the same content.
- Collective deliberation: When people over a period of time, across a variety of contexts, try to find out whether something is so, they typically assume content stability across those contexts. Consider a CIA task force concerned with whether Igor knows that Jane is a spy. They are unsure whether or not he does. Investigators, over a period of time, in different contexts study this question. If what they are trying to determine, i.e. whether Igor knows that Jane is a spy, changes across contexts, contingent, for example, on their evidence, what is contextually salient, the conversational context, etc. collective deliberation across contexts would make no sense.
- Intra-personal deliberation Suppose Igor, on his own, is trying to determine whether p is so. Suppose its being so makes a difference to his life, but he's unsure. Sometimes he thinks the evidence, on balance, supports p, sometimes not. It depends on how he looks at

the evidence, on what he takes to be the relevant considerations. Just as in the inter-personal case, this presupposes a stable content he's deliberating about.

- **Justified belief:** Much of our knowledge of the world is based on testimony. Hearing a trustworthy person assert that *p* can provide good reason to believe that *p*. If we think everything Jason says is true and he says naked mole rats are blind, we have good reason to believe naked mole rats are blind. But this is possible only if we can say what he said, viz. that naked mole rats are blind. We need to understand (and remember) what he said. We have to be in a position to agree with it. This is possible only if content can be shared across contexts.
- **Responsibility:** We hold people responsible for what they say, ask, request, claim, etc. We can do so only if we, in another context, can *understand* what they said (suggested, ordered, claimed, etc.), *say* what they said, and *investigate* what they said.
- **Reasons for actions:** A closely connected phenomenon is this: What others say often provides reasons for action. What people said in another context can provide reason for action only if we can understand what they have said, investigate it, trust it, etc.

These interconnections and mutual dependencies between content stability and non-linguistic practices are significant because any theory that implies content is *not* shared across contexts must account for the devastating implications that this view has for these non-linguistic practices. To endorse a view that implies that what we do in all these cases is based on a fundamental confusion we harbour about the nature of our own language is an awfully high price to pay to protect relevance theory.

In sum, to accept (NSC) isn't just to reject some cutesy and fluffy philosophical dogma. It has direct implications for some of the most important aspects of our non-linguistic social practices. It implies that these practices are based on fundamentally mistaken beliefs. We believe we can give the same orders to many people, that what others say can justify our beliefs, that we can hold others responsible for what they have said, etc. But there's no such thing as shared content to underwrite these practices. Of course, we could be fundamentally mistaken about ourselves in just these very ways, but at least this much is clear: If you are inclined to bite this bullet, then you had better provide an alternative account of these non-linguistic practices. Absent some such story, we

have, we think, some very good reasons for staying far away from this particular kind of revisionism.

1.3 objections to the conservative strategy

At the heart of the Conservative Strategy is the idea that sentences like ‘A said that p’, ‘A said what B said’, ‘I agree with what A said’, ‘I understand exactly what I said’, and the other such locutions do *not* require for their truth content *identity* across contexts. All that they require is content *similarity* across contexts. The details can be elucidated in various ways, one version of which is:

- ‘A said that p’ means the same as ‘A said something similar to p’.
- ‘A said what B said’ means the same as ‘A said something similar to what B said’.
- ‘A and B agree’ means the same as ‘A and B endorse similar thoughts’.
- ‘A understands what B said’ means (something like) ‘A grasped a proposition similar to the one expressed by B’.

And so on for other cases. According to this similarity view (SV), we do not make false claims when reporting or repeating others. Our practice has, wisely, factored in that there is no cross-contextual content identity.

We discuss four concerns we have about this strategy:

criticism 1: identity is transitive; similarity is not.

The most obvious and most serious problem for the Conservationists is this: Similarity is not transitive. Transitivity is, however, built into the ‘said that’ locution. Consider first (T1):

(T1): If A said what B said, and B said what C said, then A said what C said. If, however, A said something similar to what B said, and B said something similar to what C said, it simply doesn’t follow that A said something similar to what C said. (T1), according to the similarity theory, could be false. Since T1 can’t be false, the similarity theory fails. ‘A said that p’ simply does not mean the same as ‘A said something similar to p’.

The same point is brought out by (T2):

(T2): If A said that p, and B said that p, then A and B said the same (or: A said what B said).

Again, the similarity theory implies that (T2) can come out false: If A said something similar to p, and B said something similar to p it doesn't follow that what A said is *identical* to what B said (or that what B said is similar to what A said).

As far as we can tell, this is about as close to a conclusive objection to the view that 'A said that p' means the same as 'A said something similar to p' that any will ever come up with. So, we could just end the critical discussion here. But there are three other philosophically revealing objections worth mentioning.

criticism 2: sv doesn't explain our distinction between saying exactly what someone said and saying something similar but not identical?

If 'A said that p' means 'A expressed a proposition similar to p', then how do we interpret sentences like:

- He almost said that p, but didn't.
- He came very close to saying that p, but didn't.
- What he said was similar to p, but not exactly p.

The easiest way to focus this criticism is to think about (SA):

(SA): A didn't say that p, but she said something similar to p.

In uttering (SA), we don't mean what the similarity theory predicts. According to the similarity theory, 'said that' means 'said something similar to', so (SA) should mean:

(SAS): A didn't say something similar to p, but he said something similar to something similar to p.

That is not what (SA) means. Suppose A asserted q. Suppose q is similar to p. Under those conditions, SAS would be false. SA would clearly be true. So SA can't mean what SAS means.

In sum: If content similarity is employed to explain what's meant by 'saying the same', it becomes impossible to explain what's meant by 'saying something similar, but not identical'.

criticism 3: claims about degrees of similarity and comparative similarity are unintelligible in connection with 'said that' claims.

We can make intelligible and even true similarity judgements of the form:

- A is more similar to B than to C.
- A is a little bit like B.
- A is like B in some respects.

According to SV, 'A said that p' means 'A expressed a proposition that's similar to p' but that predicts we should not only be able to make sense of, but also make, true judgements of the form:

We should expect it to be possible to modify 'said' with 'very very much',

- A said p very much.

The only reading we can get of this is that he said p over and over again. We clearly cannot get the reading predicted by the similarity theory.

criticism 4: similarity without identity.

The following seems plausible: to understand what it is to say a proposition similar to p, we must understand what it is to say that p. Similarity is parasitic on identity, in at least that respect. If the similarity theory were correct, then language would leave us with no linguistic device for talking about sameness of sayings, but only with a device for talking about the derivative phenomenon. But, if, in general, talk about F's being similar to G's requires that we can talk about F's and G's directly, then wouldn't it be exceedingly surprising if our language had *no* device whatsoever, for talking about what people say rather than what's similar to what people say? (Of course, if their view is correct, there should be no way to interpret the previous sentence in the way we are sure you, our reader, as a matter of fact, interpreted it.)

1.4 nsc and evaluations of speech acts

So far we have focused on Conservatism's and Revisionism's failures to account for our reporting practices. We turn now to their implications for our *assessments* of what others say. There's an obvious connection between these two issues: to evaluate what someone has said (asserted, suggested, etc.), we have first to determine what she has said (asserted, suggested, etc.). So RT's failure to account for our practices of reporting what others have said has direct implications for our ability to assess reported contents. The way this plays out depends on whether we are considering a Conservative or Revisionist version of RT. There are three main worries:

1. *Massive verbal disagreement*: The most obvious worry is this: If (NSC) were true, then when we think we're evaluating what Naomi intended to communicate, we're really just evaluating some proposition similar to that proposition. When we think we're evaluating what Naomi said, we're not. When we think we're collecting evidence and arguments against her claim, we're not. Our evaluations are always off their mark. We never really agree or disagree – we always miss our real target. There can be no doubt that our real target is the proposition that the speaker intended to communicate – what she intended to commit herself to.

It should be obvious that neither Conservatism nor Revisionism can alleviate this problem. The Revisionist, it would seem, endorses this view. She says this is one of the things we'll see clearly once we realize RT is true; it's a view we should embrace. Again, we hope it is clear that this is no minor cost of RT. It certainly is not a view that can be endorsed in passing, as an interesting corollary. It needs an account of the point of evaluation, on the assumption that evaluations are always off target. (Indeed, if what we argue in section 2 is correct, they are typically *way off target*.)

The situation is no better for the Conservative. According to her, it is true to say that Naomi said that *p*, even though the proposition she intended to communicate is *q*. It follows that an utterance of 'What Naomi said is true' can be true, even though the proposition she intended to communicate, i.e. that *q* is false (on the assumption that *p* and *q* are R-related). This is in itself an extremely counter-intuitive and unattractive implication of (NSC).

2. *Varying truth-values*: Here's another unattractive implication of Conservatism: If the Conservative were correct, it could be true, in a context of interpretation, *C1*, to utter 'What Naomi said is true' and in another context of interpretation, *C2*, false to utter the very same sentence (assuming the demonstrative demonstrates the same utterance by Naomi). Imagine a context of utterance, *C*, in which *A* utters *S* intending to communicate *q*. In *C1*, *p* is sufficiently similar to *q*, and *q* is true, so in *C1*, it is true to say, 'What *A* said is true'. In another context of interpretation, *C2*, *p* is sufficiently similar to *r*, and *r* is false. So, in *C2* it is false to say 'What *A* said is true'. So, in *C1*, what *A* said is true, while in *C2*, what she said is false. The truth-value of her utterance will vary from one context of utterance to another.

These are two intrinsically problematic implications of (NSC). But they also lead RT into various positions best described as weakly self-defeating.

3. *NSC is communicatively self-defeating*: Hawthorne (2004) describes a position as ‘weakly self-defeating’ if the truth of the position implies that we should not believe it. (NSC) has a related problem: Let’s say a position is Communicatively Self-Defeating if it follows from the truth of the position that it can’t be communicated. If a position is communicatively self-defeating, then it implies that no one can understand it or that no one can evaluate whether it is true or false (rational or irrational, well supported or not, etc.).

(NSC) is communicatively self-defeating: it implies, first, that what the proponents of (NSC) think and intend to communicate will not be communicated to their intended audience. Some proposition R-related to what they intend to communicate might be communicated, but this might be very different from what the proponents of (NSC) have in mind. As a consequence, it implies that those who try to evaluate (NSC) (like us) will provide evidence and arguments against a position different from the position (NSC) proponents intended to communicate.

2. relevance theory and nsc

Our attribution of (NSC) to RT proponents has been based primarily on a passage in Carston (2002), where she endorses something much like it. That passage does not, however, reveal the full extent of relevance theorists’ commitment to (NSC). It does not make clear just how radical a version of (NSC) follows from taking relevance theoretic principles. In what follows we will argue:

- a. (NSC) is a direct implication of central tenets of RT.
- b. RT does not guarantee significant similarity between the proposition the speaker intended to communicate and the explicatures the audience reaches by using the procedure described by RT.

In this respect, we think, RT proponents misrepresent their theory to a significant extent. The discussion of specific examples by RT proponents essentially ignores (NSC). It gives the impression that we can say and grasp what’s said by an utterance.

2.1 optimal relevance and contextual effects

According to RT, an interpreter encounters an utterance of a sentence with a certain logical form; she then looks for a development of that logical form that is *optimally relevant*. A development of a logical form is optimally relevant just in case it is ‘at least relevant enough to warrant

the addressee's attention and, moreover, as relevant as is compatible with the communicator's competence and goals' (Carston 2001, 6–7). More specifically:

optimization of relevance

A speaker (or more generally, an ostensive communicator) calls for an expenditure of mental effort from an addressee (an outlay of attentional and inferential resources) and that licenses an expectation of a worthwhile yield of cognitive effects and no gratuitous expenditure of effort. This is captured by the 'Communicative Principle of Relevance': every act of ostension communicates a presumption of its own optimal relevance; that is, a presumption that it will be at least relevant enough to warrant the addressee's attention and, moreover, as relevant as is compatible with the communicator's competence and goals. (Carston 2001, 6–7)

The presumption of optimal relevance leads us to utilize the least effort strategy (LES) for interpretation (which we repeat here):

(LES) Check interpretive hypotheses in order of their accessibility, that is, follow a path of least effort, until an interpretation which satisfies the expectation of relevance is found; then stop. (Carston 2001, 6)

Carston spells out the reasoning in the following passage:

The least effort strategy follows from the presumption of optimal relevance in that the speaker is expected to have found an utterance for the communication of her thoughts which minimizes the hearer's effort (modulo her own goals and abilities); the justification for the addressee stopping processing as soon as an interpretation satisfies his expectation of relevance follows similarly, in that any other interpretation that might also achieve the requisite level of effects will be less accessible and so incur greater processing costs. (Carston 2001, 6–7)

To understand this, we need to understand the central notion in RT: *Relevance*. The notion of relevance is intertwined with that of a *Contextual Effect*. Sperber and Wilson write:

The notion of a contextual effect is essential to a characterization of relevance. We want to argue that having contextual effects is a

necessary condition for relevance, and that other things being equal, the greater the contextual effects, the greater the relevance. (Sperber and Wilson 1986, 116)

The connections between relevance and cognitive effects are spelled out in the following two conditions:

Extent Condition 1: An assumption is relevant in a context to the extent that its contextual effects in this context are large.

Extent Condition 2: An assumption is relevant in a context to the extent that the effort required to process it in this context is small. (Sperber and Wilson 1986, p.125)

Carston summarizes the notion of a contextual (or cognitive) effect, as follows:

Cognitive effects (or contextual effects) include the strengthening of existing assumptions of the system, by providing further evidence for them, the elimination of assumptions that appear to be false, in the light of the new evidence, and the derivation of new assumptions through the interaction of the new information with existing assumptions. (Carston 2001, 6)

On this view, a new assumption (or a new piece of information) can have three kinds of contextual effects: (a) it can *combine* with old information to derive new information (that couldn't have been derived before); (b) it can provide new evidence *for* old assumptions; and (c) it can provide new evidence *against* old assumptions. An interpretation is relevant to the extent that it has contextual (or cognitive) effects.

2.2 from optimal relevance to nsc

When an interpreter looks for an interpretation that satisfies the principle of optimal relevance, she checks for cognitive (i.e. contextual) effects. The cognitive effects of an utterance on a person at a time *t* will depend, essentially, on the beliefs the interpreter has at *t*. These vary between interpreters. They vary for a single interpreter over time. Here is an example: The sentences of this paper have certain logical forms. The readers of this paper will develop these until they satisfy the Principle of Optimal Relevance. Which development satisfies that principle for a particular reader *R* will depend on the contextual effects these logical forms have on *R*. We have no way to predict in advance

which development of these logical forms various readers will end up with. There are infinitely many such developments and common sense dictates that readers will all end up in different places. There's not even a guarantee that these places will be similar – developments of these logical forms can be radically different and it would be a minor miracle if they were not.

When Carston and Bezuidenhout tell us to relinquish the very idea of shared content, they proffer the prospect of similarity as a substitute, but what does that really amount to and what in RT guarantees it? Any interpretation of an utterance *u* of a sentence *S* will be a development of the logical form of *S* – so in that sense, there's similarity. But what counts as a 'development' is unconstrained in this sense: whatever you end up with as a result of applying (LES), starting from the logical form of the sentence you're interpreting, will be a development of that logical form. What you end up with depends on what cognitive effects the utterance has on you, and that again on what beliefs (etc.) that the interpreter has. Even if we fix a standard of similarity, there's no guarantee that what radically different interpreters would end up with is similar were they to use (LES).

The way we just put this point assumes that there's some fixed standard of similarity that RT can appeal to. That assumption is false. It's a basic fact about similarity that what counts as similar to what depends on contextually determined standards of similarity. So if the relevance theorist tells us that the process she describes ends up guaranteeing similarity between the results reached by different conversational participants, she needs to tell us *by what standard of similarity* this is supposed to be so. There is trivially no guarantee that what we end up with will be similar *by all such standards*. Of course, there will always be *some* standards by which what interpreters and speakers end up with is similar, but that's hardly comforting (since this standard could be absurdly loose – e.g. explicatures will all be similar in that they are developments of logical forms). The relevance theorist could try to suggest that the interpretative results in *C* will be similar by the similarity standards of *C*. But again, given the radical variability in standards between contexts, what could possibly guarantee this claim? Certainly nothing in Relevance Theory itself guarantees it.

In sum, the claim that what conversational participants end up with will be similar, has no basis in relevance theory. For a most radical illustration of this, consider Carston's notion of *ad hoc concept construction*; these are 'cases where any one of a wide range of related concepts might

be communicated by a single lexical item' (Carston 2004, 641). She gives the following examples:

... think of all the different kinds, degrees and qualities of feeling that can be communicated by each of *tired, anxious, frightened, depressed, well, happy, satisfied, sweet*, etc. 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, 644)

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. For the reasons given above, it is not only hard to imagine an argument that could establish that a speaker and number of audience members will invariably end up with similar results after going through such a process, it also hard to understand what 'similar' could mean in such a claim.

In conclusion: Proponents of RT should endorse a particularly extreme version of NSC, according to which speaker and audiences might faultlessly interpret each other, and nonetheless end up with explicatures that are radically different from each other.

3. pluralistic minimalism

A central challenge in pragmatics is to develop a theory of communication that reconciles two fundamental facts: we can share contents across contexts and communicated content is deeply context sensitive. Both sets of data are robust: what ends up in a speaker's mind when hearing an utterance depends on an extraordinarily wide range of features both of the speaker's context of utterance and the context of the audience. At the same time, the idea that speakers and audiences, variously situated, should be incapable of sharing contents runs counter to equally

fundamental features of our linguistic self-understanding (if that sounds too metaphorical, what we mean by it is spelled out in section 1).

Here, in summary form, is our solution to this apparently paradoxical data set (presented in more detail in Cappelen and Lepore 2004): The crucial step is to relinquish what we call Speech Act Monism. This is the view that each utterance of a sentence says (asserts, claims, etc.) just one thing (one proposition, one thought). It is Speech Act Monism that generates even the appearance of tension between content stability and variability: If utterance u_1 of S says just one thing, e.g. p , and utterance u_2 of S says something else, e.g. q , and if $p \neq q$, then how could u_1 and u_2 say the same? Here's the solution: Drop the idea that an utterance expresses one proposition, i.e. endorse a combination of what we call Speech Act Pluralism and Semantic Minimalism. (We call the combination Pluralistic Minimalism.)

According to Speech Act Pluralism, any utterance can be used to express a plurality of propositions. Accordingly, u_1 of S expresses a set of propositions, say, $C1$, and u_2 of S expresses a set of propositions, say, $C2$; and it may be that $C1 \neq C2$, i.e. they don't share the exact same members. This, however, does not prevent an overlap. If $C1$ and $C2$ do overlap, then there is an obvious explanation of how u_1 and u_2 can both say different things and yet say the same. When we speak of two utterances of S saying the same, we are focusing on the area of overlap, and when we speak of two utterances saying different things, we are focusing on the area of non-overlap.

Two utterances of S might express different sets of propositions. We claim that if you adjust for obviously context sensitive expressions (i.e. hold the semantic value of these stable), then these sets will have at least one proposition in common. Call this *the semantic content of S* , i.e. one way (not the only) to characterize the semantic content of S is as that content which all utterances of S have in common (once we adjust for obvious context sensitivity). The view that there is such a common content we call *Semantic Minimalism* ('minimalism' because the contextual influence is minimal).

What's our argument for Semantic Minimalism, i.e. that there is such an overlap between different utterances of S ? In earlier work, we presented three kinds of arguments (for more details, see Cappelen and Lepore 2004, ch. 10):

1. Semantic Minimalism helps explain how we can share contents across contexts. If we accept that theory, we can explain why contents are not contextually trapped. If our arguments above are right, then this

is our only protection against what can be called contextual content solipsism. Semantic Minimalism guarantees a level of content that enables speakers whose conversational, perceptual and cognitive environments are very different to agree and disagree. It's a kind of inference to the best explanation.

2. There's a related argument (in some sense the flip side of the last one), but it appeals more directly to intuitions: When we encounter a range of utterances of S in diverse contexts (or just one utterance in a context we are ignorant of), we're often inclined to use S to say what was said by these utterances. When we do that, i.e. when we focus on what they all share (or what was said by a single utterance in an unknown context), we have a kind of direct access to the minimal content. It's not something we focus on (or care about) in most contexts, but when we do, it's right there and we have direct cognitive access to it.
3. Finally, we argued that the view that there's no common content is internally inconsistent. We will not present that argument here because it requires saying much more about our opponent's position, but for an extended discussion see Cappelen and Leopre 9 (2005, ch. 9).

Pluralistic Minimalism has two important corollaries:

1. A Pluralistic Minimalist must reject the Speech Act Conception of Semantics. The Speech Act Conception of Semantics can be characterized as follows:
 - According to the Speech Act Conception of Semantics variability in what speakers say by uttering S in different contexts is relevant to semantics of S because the goal of semantics is, roughly speaking, to account for the content of speech acts performed by utterances of sentences. So, if S is a sentence of L and S is used to say that p (to assert that p), then the semantics for L should explain how that could be. On this view, there must be a close explanatory connection (this connection can be spelled out in various ways) between the semantic content of S and the content of speech acts involving S. As a corollary, if what is said by utterances of S varies between contexts of utterance, then the semantic content of S should be context sensitive.
2. Pluralistic Minimalists must also reject the Semantic Conception of Indirect Reports, according to which If 'A said that p' is a true indirect report of an utterance of S, then the semantic content of p (as it occurs in that report) should be identical to the semantic content of S. In short, they reject that view that indirect reports report on semantic contents.

Both (1) and (2) have important implications for how one thinks about the methodology of semantics, (for more details, see Cappelen and Lepore 2004, ch. 1–4). Even with these pieces in place, we should emphasize that Pluralistic Minimalism is a research project, not a completed theory. The challenges for this projects fall into three main categories:

- Say more about the nature of speech act content.
- Say more about semantic contents.
- Say more about the relationship between speech act content and semantic content.

We have started some of that work in Cappelen and Lepore (2006) and (forthcoming), but we have no doubt that the framework is in need of considerable additional refinement.

references

- Bezuidenhout, A. (1997). 'The Communication of *De Re* Thoughts'. *Nous* 31.2: 197–225.
- (2002). 'Truth-Conditional Pragmatics'. *Philosophical Perspectives* 16: 105–34.
- Cappelen, H. and E. Lepore (2004). 'Context Shifting Arguments'. *Philosophical Perspectives*. 25–50.
- (2005). *Insensitive Semantics*. Oxford: Basil Blackwell.
- (2006). 'Reply to Critics'. *Mind and Language*.
- (forthcoming). 'Reply to Critics'. *Philosophy and Phenomenological Research*.
- Carston, R. (2001). 'Explicature and Semantics'. In S. Davis and B. Gillon (eds). *Semantics: A Reader*. Oxford: Oxford University Press.
- (2002). *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Oxford: Blackwell.
- (2004). 'Relevance theory and the saying / implicating distinction'. in L. Horn and G. Ward (eds). *The Handbook of Pragmatics*. Oxford: Blackwell 633–56.
- Hawthorne, John (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Recanati, F. (2001). 'What is Said'. *Synthese* 128: 75–91.
- (2004). *Literal Meaning*. Cambridge: Cambridge University Press.
- Searle, John (1978). 'Literal Meaning'. *Erkenntnis* 13: 207–24.
- (1980). 'The Background of Meaning'. In J. Searle, F. Kiefer, and M. Bierwisch (eds). *Speech Act Theory and Pragmatics*. Dordrecht: Reidel. 221–32.
- Sperber, D. and D. Wilson (1986). *Relevance*. Oxford: Blackwell.
- Travis, Charles (1985). 'On What is Strictly Speaking True'. *Canadian Journal of Philosophy* 15.2: 187–229.
- (1989). *The Uses of Sense*. Oxford: Oxford University Press.
- (1996). 'Meaning's Role in Truth'. *Mind* 100. 451–66.