

Relevance Theory and the Problem of Content Sharing

Master thesis in Philosophy

Jan-Fredrik Braseth



Department of Philosophy, Classics, History of Arts and Ideas

Supervisor: Carsten M. Hansen

UNIVERSITETET I OSLO

May 2010

Acknowledgements

First and foremost, I would like to thank my supervisor, Carsten M. Hansen, for helping me choose the topic of this thesis. He has also provided me with valuable guidance during the whole process, which I am very grateful for. I also want to thank Georg Kjøll for some helpful discussions about topics related to this thesis. Finally, I want to thank Gudrun Breivik for helping me with proof reading, and for her general support during the last part of my process.

Table of Contents

Abstract.....	5
1 Introduction.....	6
2 Relevance Theory – An overview.....	9
2.1 Manifestness.....	10
2.2 The communicative intention.....	11
2.3 Understanding an utterance.....	12
2.3.1 The Cognitive Principle of Relevance.....	12
2.3.2 The Communicative Principle of Relevance.....	15
2.3.3 Getting the propositional form.....	16
2.3.4 The Least Effort Strategy.....	19
2.3.5 Identifying the implicatures.....	20
3 A presentation of Cappelen and Lepore's arguments.....	22
3.1 The Non-shared content principle and Relevance Theory.....	22
3.2 The consequences.....	26
3.2.1 Appeal to similarity.....	28
4 What Relevance Theory is committed to.....	29
4.1 What a speaker wants to communicate.....	29
4.2 What the hearer thinks he has grasped.....	31
4.3 The revised non-shared content principle.....	32
4.4 Duplication of thoughts.....	32
5 What Relevance Theory should be committed to.....	35
5.1 Predicting what hearers end up with.....	35
5.1.1 Constraints given by the Logical form.....	35
5.1.2 Non-linguistic constraints.....	38
5.1.3 Mutually manifest facts.....	38
5.1.4 Contextual constraints.....	39
5.1.5 Constraint given by the Communicative Principle of Relevance.....	40
5.1.6 The constraints cannot guarantee success.....	41
5.2 The unarticulated argument.....	41
5.3 Content.....	43

	4
5.3.1 Cappelen & Lepore's use of 'content'.....	43
5.3.2 Defining content.....	45
5.3.3 Revised version of the unarticulated argument.....	48
5.3.4 The difference between fixed filbert content and explicated content.....	49
5.4 Why is propositional fixed filbert content necessary?.....	51
5.4.1 The need for a failsafe procedure.....	51
5.4.2 C&L's minimal proposition does not function as a safety mechanism.....	53
5.5 Where do we stand?.....	56
6 Communication without content sharing.....	58
6.1 Resemblance.....	58
6.2 What is the aim of communication.....	60
6.2.1 Successful communication.....	62
6.3 Resemblance does not lead to the Similarity View.....	66
6.4 Examples of successful communication.....	67
7 Summary and conclusions.....	71
8 Bibliography.....	73

Abstract

When a speaker communicates with someone, she wants to convey some kind of content to the hearer. For the communicational act to be successful, the hearer will have to grasp this content. Human communication seems, in this way, to rely on content sharing. Relevance Theory is a pragmatic theory aiming at giving an account of the mechanisms underlying human communication. Herman Cappelen and Ernie Lepore have argued that Relevance Theory implies that content sharing is impossible. This thesis is an attempt to defend Relevance Theory against this claim.

I start by arguing that Cappelen and Lepore's argument is imprecise, and suggest a modification of their claim. Specifically, I make some adjustments to what they call 'The Non-Shared Content Principle'. According to Cappelen and Lepore, it is the commitment to this principle that makes Relevance Theory imply that content sharing is impossible.

I then show that Cappelen and Lepore's argument relies on the idea that the proposition expressed must be given by the semantics of the language alone, without any intrusion from the context. I argue that this is impossible; The hearer needs *some* contextual guidance to be able to grasp the proposition expressed. With that in mind, I argue that even though the proposition expressed is not determined by semantics alone, this does not mean that a hearer will be unable to grasp it.

Finally, I argue that one does not necessarily need to grasp the same proposition for communication (and content sharing) to be possible. In most cases, people can understand each other without grasping the same proposition, as long as they grasp propositions that are *similar* to a certain extent. I argue that this is because communication is not about conveying one specific proposition, but about conveying a 'point'.

1 Introduction

Quite often, it seems, there is a divergence between what people say and what they mean. The sentence a speaker uses to express herself does not always mean the same thing as what she means. The most obvious example is when someone is being ironic. In these cases it seems as though the speaker means *the opposite* of what she says. In cases of metaphor the sentence will sometimes seem to say something trivially false. A person might say to a loved one: “You are my anchor in the storm”, and even though the hearer is not an anchor, he would probably still seem to think that what was expressed not only was something true, but something meaningful as well.

In other cases the sentence might reveal part of what the speaker means, but the speaker's intention is to use this to *imply* something else. Someone might say that “it is cold”, and by that mean that they want the hearer to close the window. Most of the time, the hearer would, as we say, “take the hint” and close the window. Sometimes the utterance is not even a complete sentence; for example when someone says “higher”, and by that is conveying that they think a certain picture should be hung higher on the wall. Even though the speaker utters only one single word, the hearer will probably understand what she means.

All these cases fall under what has been known as the *Underdeterminacy Thesis* (Carston 2002, p. 19). In some ways, the linguistic meaning of an utterance underdetermines what is actually said or meant by the speaker¹. Any theory of language use should be able to explain how the hearer can understand what has been meant by the speaker, even given, what seems to be, lack of evidence. The fact is that even when the sentence used does not express the speaker's meaning, hearers, very often, successfully understand what the

¹ Actually there are three slightly different variations on the Underdeterminacy Thesis. They depend on exactly what one means by the terms 'what is said' and 'what is meant' (Carston 2002, p. 19). However, the general point is that the linguistic meaning of an utterance is, in itself, not enough to determine what the speaker wants to communicate.

speaker means.

Relevance Theory (henceforth RT) is a pragmatic theory of human communication developed in the 1980s by Dan Sperber and Deirdre Wilson, and since the release of *Relevance* (1986) their theory has been highly influential both in the field of linguistics and philosophy of language. Its main goal is to give an account of the mechanisms underlying human communication, by attempting to give an answer to how hearers go from an underdetermined utterance to an idea about what the speaker actually means. The central claim of RT is that an utterance raises certain specific, and predictable expectations of *relevance*². These expectations guide the hearer towards the speaker's meaning (Wilson/Sperber 2004, p. 607).

RT is a cognitive psychological theory, and like most other psychological theories, it has testable consequences that can be confirmed, disconfirmed or fine-tuned in the light of experimental evidence (Wilson/Sperber 2004, p. 625)³. Being a study of language use, it is also a pragmatic theory. Pragmatics is an empirical science, but it has “philosophical origins and philosophical import” (Sperber/Wilson 2005, p. 468). This thesis will be a discussion of the more philosophical imports of the theory rather than the empirical ones.

Cappelen & Lepore (henceforth C&L) have, in several places, argued that the account RT gives is, in some ways, lacking (Cappelen/Lepore 2005; 2007; 2008). They argue that if RT were right in its description of how a hearer grasps the speaker's meaning, people would almost never understand each other. RT implies, according to C&L, that 'content sharing' is impossible (Cappelen/Lepore 2007, p. 1). A speaker may want to communicate some specific content, but she will never get this content across to the hearer. This view is what C&L call *The Non-Shared Content Principle* (Cappelen/Lepore 2007, p. 3). They claim that

2 In RT, 'relevance' is a technical term and does not necessarily mean the same thing as what is normally meant when using the term. I will define the technical notion of 'relevance' in section 2.3.

3 For more information on what sort of experiments one can devise to confirm RT, see (Sperber/Wilson 2004, section 6).

if this view were true, it would have major consequences in our everyday lives, and see this as an argument for the falsity of RT; Since it seems as though we *are* able to share content when we communicate, the account given by RT cannot be true.

In this thesis I will investigate C&L's argument and, ultimately, defend RT against it. I will also make a case for the possibility of communicating without sharing propositional content. In the process I will present the debate as it currently stands, and clear up several misconceptions related to this debate. I will, hopefully, be able to give a clearer view of exactly what C&L's argument amounts to, what their view is, and why RT does not suffer the consequences presented by C&L.

In section 2 I will give a brief presentation of RT and its central claims. In section 3 I will present C&L's argument that RT implies the impossibility of content sharing, and the consequences this will have. In section 4 I will argue that RT is not committed to C&L's *non-shared content principle*. I will argue that this is, first of all, because of the phrasing of the principle, which is based on a slight misconception of how RT explains communication. Then I will propose a revised principle which is in the spirit of C&L's original one, but provides a more substantial challenge to RT. After that I will uncover what I call *the unarticulated argument* in C&L's argumentation, namely that semantic propositional content is necessary for communication. In section 5 I will first investigate exactly what is meant by 'content', and show that this notion is not always clearly defined and therefore leads to misunderstandings in the debate. I will then try to define several notions of 'content' in order to make the debate clearer. When this is done, I will look at C&L's argument for why semantic propositional content is necessary for communication. I will argue that this is false, and that RT can, in a lot of cases, explain how people can share propositional content even without it being semantically determined. I will show that there are, however, reasons to believe that in a lot of cases people do *not* share propositional content. In section 6 I will try to show how RT can explain communicational success even when the speaker and hearer do not grasp the same proposition.

2 Relevance Theory – An overview

In this section I will give a brief introduction to RT and its central tenets. My aim is to give a presentation that is complete enough to make the reader able to follow C&L's argument in section 3. For a fuller account of RT, see Sperber & Wilson (1995), Carston (2002) and Wilson & Sperber (2004).

RT is in many ways an expansion of Grice's attempt at an *inferential model* for communication. On his model communicating is about inferring the speakers meaning. This is opposed to the more traditional semiotic view that communication is based on a *code model* (Sperber/Wilson 1995, p. 2). The code model basically explains communication by telling a story about how a speaker having a thought, *encodes* that thought into words, and how the hearer *decodes* the words into a thought equal to the thought the speaker had in mind. The inferential model explains communication by saying that hearers infer the speakers intentions when they attend to an utterance. For more on these two models and why they, according to Sperber & Wilson, fail to give a correct model of human communication, see chapter 1 of *Relevance* (Sperber/Wilson 1995).

RT concedes that there is some kind of linguistic meaning which can be decoded from the utterance, but this meaning falls short of encoding what the speaker means:

[H]uman intentional communication is never a mere matter of coding and decoding. The fact is that human external languages do not encode the kind of information that humans are interested in communicating. Linguistically encoded semantic representations are abstract mental structures which must be inferentially enriched before they can be taken to represent anything of interest. (Sperber/Wilson 1995, p. 174)

In my account of RT, I will take a look at how it combines the two distinct processes of decoding and inference into a fuller account of what happens when people communicate verbally.

My account of RT will begin with an overview of the kind of information that RT claims is conveyed in human communication; the kind of information that a speaker wants to share with her audience when she is speaking to them.

2.1 *Manifestness*

All humans live in the same shared environment and we cannot help but make *assumptions* about this environment. Assumptions are in RT defined as “thoughts treated by the individual as representations of the actual world” (Sperber/Wilson 1995, p. 2). So, if I see a cup of coffee on the table in front of me, I will form an assumption that there is a cup of coffee on the table; If I feel cold I will entertain the assumption that I am cold.

We are, to a varying degree, capable of recognizing the physical world we live in and how it is organized. Sperber & Wilson explains this by appealing to their notion of *manifest facts*⁴:

A fact is manifest to an individual at a given time if and only if he is capable at that time of representing it mentally and accepting its representation as true or probably true.
(Sperber/Wilson 1995, p. 39)

Manifest facts can be said to be all the facts in an individual's environment that she is capable of becoming aware of. Whether or not she *will* become aware of them is dependent on her behavior and cognitive abilities. That there is a cup of coffee on the table is a manifest fact to me, but to form the assumption that there is a cup of coffee on the table I will have to turn my head in that direction, focus my attention on the cup, etc.

Sperber & Wilson extend the notion of manifest facts to manifest assumptions: “An assumption, then, is manifest in a cognitive environment if the environment provides sufficient evidence for its adoption” (Sperber/Wilson 1995, 39). They also insist that their point of view is “cognitive rather than epistemological” (Sperber/Wilson 1995, p. 39). What

4 Note that the term 'manifest' is a technical term in RT. It does not necessarily correspond to the normal definition as something that is 'apparent' or 'obvious'. See the discussion on how manifestness comes in degrees further down.

is meant by this is that assumptions do not necessarily correspond with how the world actually is. Any assumption, whether true or false, may be manifest to an individual. From a cognitive point of view “mistaken assumptions can be indistinguishable from genuine factual knowledge” (Sperber/Wilson 1995, p. 39).

Sperber & Wilson also say that there are degrees of manifestness; The assumptions that are more likely to be entertained by an individual are also more manifest to her. The individual's cognitive abilities and the state of the physical environment around her will decide to what degree an assumption is manifest to her (Sperber/Wilson 1995, p. 39). It is more manifest to me that there is a cup of coffee on the table, than that there is a piece of gum stuck under the table. It will require some more work for me to form the latter assumption (I will have to bend down and look under the table).

Facts can also be *mutually manifest* to a group of people. A fact is mutually manifest to a group of people, if it is manifest to each individual in the group that a fact is manifest to all the members of the group. For example: If Gudrun joins me at the table with the cup of coffee described earlier, the fact that there is a cup of coffee on the table will be manifest to *her* as well. In addition, the fact that this fact is manifest to both of us, is *itself* manifest to both of us. The fact that there is a cup of coffee on the table is therefore mutually manifest to the two of us.

2.2 The communicative intention

The set of manifest facts for an individual is called her *cognitive environment*, and when we are communicating, our goal is, according to RT, to modify each other's cognitive environment. According to RT the communicator has an *Informative intention*. This intention is “to make manifest or more manifest to the audience a set of assumptions **I**” (Sperber/Wilson 1995, p. 58). If I am telling Gudrun that there is coffee on the table, one of my goals is to make it manifest to her that there is coffee on the table (in case she has not

noticed it). Other goals might be to make it manifest to her that I am offering her coffee, that she is free to have some coffee, that I have made coffee, etc.

Further, for it to be classified as communication, the hearer needs to be aware of the fact that the communicator has this informative intention. If I just make a noise, Gudrun might turn her head in my direction, see the cup of coffee on the table and, as a result, form the assumption that there is coffee on the table. Even if I, for some reason, made the noise for just that purpose, I would not be *communicating* to Gudrun that there is coffee on the table. For it to be communication, Gudrun must know that I want to communicate something. RT says that the communicator, in addition to having an informative intention, also needs to have a *Communicative intention*. Its goal is “to make it mutually manifest to audience and communicator that the communicator has this informative intention” (Sperber/Wilson 1995, p. 61). To put it simply: When someone wants to communicate something, she must make sure that the audience are aware of her intention to communicate with them.

2.3 Understanding an utterance

According to RT then, when we communicate verbally our goal is to use language to make a set of our assumptions **I**, manifest to our audience. But how, one might wonder, is the hearer able to recover those assumptions, given the underdeterminacy of language? In this section I will briefly sketch RT's explanation of the mechanisms involved in the recovery of speaker meaning. First of all, there is a need to explain why we even start paying attention when someone utters a sentence. The answer to this lies with the two principles of relevance.

2.3.1 The Cognitive Principle of Relevance

The *Cognitive Principle of Relevance* is the principle that: “[h]uman cognition tends to be geared to the maximisation of relevance” (Wilson/Sperber 2004, p. 610). RT claims that humans always pay attention to the most relevant input in a certain situation. This is not a

conscious choice, but happens automatically. RT claims that this is because our cognitive apparatus has evolved to attend to whatever is most relevant to us. There are always a lot of things happening around us, and our mind is only able to focus on a few of them at the time. To increase efficiency our cognitive system has developed in a way that it tries to pick out whatever is more relevant in a given context (Wilson/Sperber 2004, p. 610). We may be focused on making dinner, but if we hear the sound of glass breaking behind us, we automatically turn to attend to this new information. What determines whether something is relevant in a given context? To answer this, an explanation of the term *relevance* is needed.

In RT 'relevance' is a technical term, so it does not necessarily correspond to the way we normally use it. The notion of 'relevance' is closely tied with the notion of *cognitive (or contextual) effect*⁵. A cognitive effect is an effect that makes the individual change some of her assumptions. It can either be a strengthening or a weakening of existing assumptions, or the acquiring of new assumptions (Carston 2001, p. 6). Furthermore, a *positive cognitive effect* is defined as a cognitive effect that makes a worthwhile difference to the individual's representation of the world. A true conclusion will be a positive cognitive effect, whereas a false conclusion will be a cognitive effect but not a positive one, since false conclusions are not worth having (Wilson/Sperber 2004, p. 608). In relevance-theoretic terms, an input is relevant to an individual "when its processing in a context of available assumptions yields a positive cognitive effect" (Wilson/Sperber 2004, p. 608).

However, relevance is a matter of degree. Something can be more or less relevant in a given context. The degree of relevance is decided by two factors: the magnitude of the positive cognitive effect, and the magnitude of the processing effort required to achieve this effect. The magnitude of a cognitive effect can intuitively be said to be decided by how worthwhile the conclusions to an input are (Wilson/Sperber 2004, p. 609). Noticing that my train is two minutes late, might give me a small positive cognitive effect. Noticing that my train is half an hour late, however, may lead me to reorganize my appointments that day,

⁵ Originally Sperber and Wilson used the term *contextual effect*. In the post face of *Relevance* (1995) they define *cognitive effect* as "contextual effects in an individual" (p. 265). These terms seems to sometimes be interchanged (E.G. Carston 2001).

etc. Noticing that the train is half an hour late is therefore a more 'worthwhile' conclusion, than that the train is only two minutes late. In the same way, the amount of processing effort needed to process a stimulus will determine its relevance:

[T]he greater the effort of perception, memory and inference required [to process a certain input], the less rewarding the input will be to process, and hence the less deserving to our attention. (Wilson/Sperber 2004, p. 9).

A definition of relevance can thus be given:

Relevance of an input to an individual

- a. Other things being equal, the greater the positive cognitive effects achieved by processing an input, the greater the relevance of the input to the individual at that time
- b. Other things being equal, the greater the processing effort expended, the lower the relevance of the input to the individual at that time. (Wilson/Sperber 2004, p. 9)

It is worth noting that, although this is a comparative definition, it doesn't always let us decide which of two inputs are most relevant, since it relies on two separate conditions. This can be illustrated by an example (taken from Wilson/Sperber (2004, p. 609)); Mary, who dislikes most meat and is allergic to chicken, rings her dinner party host to find out what is on the menu. He could *truly* tell her any of three things:

- (1) We are serving meat
- (2) We are serving chicken
- (3) Either we are serving chicken or (7²-3) is not 46

All the utterances would be relevant to Mary, but (2) is the most relevant one. By comparing (1) and (2) one can see that (2) is more relevant for reasons of cognitive effect (and no, or a very small amount of, difference in processing effort). Comparing (2) and (3) will tell us that (2) is more relevant, because they both have the exact same amount of cognitive effect (because they are logically equivalent), but (3) requires more processing effort. However it is not straightforward to tell which is more relevant of (1) and (3). This is because there is a trade-off between cognitive effect and processing effort. If Mary is only a little bit allergic to chicken then (1) might be more relevant. If she is heavily allergic and might die if she tastes chicken, then the processing effort in (3) might be worth it, and thus make (3) more relevant.

2.3.2 The Communicative Principle of Relevance

This leads us to the second principle of relevance, called the *Communicative Principle of Relevance*. This principle states: “[E]very act of ostensive communication communicates a presumption of its own optimal relevance” (Sperber/Wilson 1995, p. 266). What they mean by the *Presumption of optimal relevance* is further defined (and then revised) in the second edition of *Relevance* (1995):

Presumption of optimal relevance (revised)

- (a) The ostensive stimulus is relevant enough for it to be worth the addressee's effort to process it.
- (b) The ostensive stimulus is the most relevant one compatible with the communicator's abilities and preferences. (Sperber/Wilson 1995, p. 270)

So when someone communicates something, the act of communicating is in itself presuming that it is optimally relevant to the hearer. The hearer will then be justified in focusing his attention to the speaker. Not only does the presumption of relevance convey that it is worth the listener's effort to attend to the utterance, but he can also be sure that the utterance is the most relevant one the speaker could produce (to a certain extent⁶).

Sperber & Wilson say that the Communicative Principle of Relevance is not a normative, but a descriptive, claim. It is not something the communicator (nor the audience) needs to be aware of, since it automatically follows with every act of ostensive communication (Sperber/Wilson 1995, p. 271). The principle is only a description of the nature of communication. It states that for communication to be possible, the hearer has to pay attention to the communicator. If the Cognitive Principle of Relevance is right, people automatically attend to the most relevant information at a given time. This means that the success of communication depends on the hearer taking the utterance to be relevant enough to be worthy of attention:

⁶ The part in the definition that says: “the most relevant one compatible with the communicator's abilities and preferences” is there to make sure that it can account for occasions where the speaker is not able to be totally relevant (lack in ability) or does not want to be totally relevant (by preference).

Wanting her communication to succeed, the communicator, by the very act of communicating, indicates that she wants the audience to see her utterance as relevant, and this is what the Communicative Principle of Relevance states. (Sperber/Wilson 2005, p. 474)

It should now be clear how the two principles of relevance make sure that the hearer pays attention to what is being communicated. The next thing in need of an explanation is how the hearer processes the information he receives.

2.3.3 Getting the propositional form

The first step in the process of understanding is a decoding process. This process is performed by an *autonomous linguistic system*, and takes the audio/visual linguistic stimuli as input and gives a semantic representation of this as output. This semantic representation is referred to as *logical form* (henceforth LF) or the *encoded meaning* of the utterance (Carston 2002, pp. 57-58). One can see the LF as a kind of 'schema' or a formula for how to construct a certain meaning from the input given. It contains various slots that have to be filled. Sometimes the filling of these slots will be partly constrained by a procedure for how to fill them (Carston 2002, p. 57). Because of the nature of the LF, it is rarely propositional⁷. There is a need for several other processes to form the LF into a fully fledged proposition, called *the propositional form* of the utterance.

Transforming the LF into a propositional form is a complex procedure involving a lot of sub-procedures. This process is being referred to as the *development* of a logical form (Sperber/Wilson 1995, p. 181). The hearer will have to disambiguate the sentence, assign referents to each referring expression and *enrich* vague terms. I will explain these processes with an example from *Relevance* (1995, pp. 177-179). Mary wants Peter to come and eat dinner, and utters the sentence:

(4) It will get cold

The logical form of this sentence is not propositional. To get to the propositional form of the sentence, Peter will have to disambiguate the term 'cold' (meaning either experiencing

⁷ Propositionality is defined in *Relevance* as “capable of being true or false” (Sperber/Wilson 1995: p. 72).

cold or inducing cold). He will also have to assign a referent to the term 'it'. Finally he will have to enrich the term 'will' (defining a time span). Doing this, he might end up with a proposition similar to this:

(5) The dinner will get cold very soon

In this case the propositional form is also one of the *explicatures* of the utterance. An explicature is a communicated assumption which is reached by developing the logical form of the utterance. This is opposed to *implicatures*, which are all the communicated assumptions that are not explicit (they can not be reached by development of the LF alone) (Sperber/Wilson 1995, p. 182). A possible implicature in this case might be:

(6) Peter should come and eat dinner at once

In addition, there is sometimes a need for a process, similar to disambiguation, called *ad hoc meaning construction*.

In some cases, words might take a slightly different meaning in the given utterance than what they normally have, and the interpreter has to infer this 'new' meaning. This can actually be seen as a general rule: “The decoded senses of a word or other linguistic expression in an utterance provide a point of departure for an inferential process of meaning construction” (Sperber/Wilson 2008, p. 181). These concepts are pragmatically constructed “on the fly” in a certain context and are not linguistically given (Carston 2002, p. 322). An example of this is the concept HAPPY which is a very general term that covers a wide range of states and feelings. In a specific context the word 'happy' might be understood as HAPPY*⁸, covering just a subset of those states and feelings.

Ad hoc concepts might be narrower or broader than the respective lexicalized concept. In the utterance:

(7) I have a temperature

uttered by someone who's feeling ill, the word 'temperature' would be taken to mean “a

⁸ I am here following the convention of writing pragmatically derived concepts with an asterisk. This is to separate it from the lexicalized concept.

temperature above normal” and hence have a meaning which is narrower than the standard lexicalized meaning. In the utterance:

(8) France is hexagonal

'hexagonal' would be taken to have a meaning which is broader than the standard meaning (since the shape of France, strictly speaking, is not hexagonal). There are also ad hoc concepts that are not reached by just narrowing or broadening the lexical meaning. Various cases of *loose use* can be shown, as in:

(9) For luggage, pink is the new black

Here 'black' is used to denote a category of fashionable colors, but the process for reaching this concept is the same as in other ad hoc constructions (Sperber/Wilson 2008, p. 188).

It should now be sufficiently clear what processes a hearer goes through to develop the logical form into a propositional one. Still, one might question how the hearer ends up with the correct propositional form. When the hearer disambiguates a term, why does he choose one, rather than another lexical meaning of the term? When hearing (4), how does Peter know that 'it' refers to the dinner? There is a need for an explanation of what guides the hearer to construct the *right* propositional form.

Another important thing to notice is that the propositional form might only be *part* of what the speaker wants to convey⁹. She might also have in mind one, or several, implicatures. In the above example, Mary did not only intend Peter to recover the explicature(5) but an implicature (6) as well. The recovery of implicatures are guided by the same principle that guides the recovery of the propositional form. This guiding principle will be presented in the next section.

9 In some cases it is arguably *not* something the speaker wants to convey. For example in cases of irony.

2.3.4 The Least Effort Strategy

We are in need of an explanation of how the hearer is able to interpret an utterance in the correct way. Why does he end up with the explicatures and implicatures that he does, when it seems to be many different interpretations of the same sentence? The reason is that the interpreter is not going through this process blindly, but is (automatically) following a specific strategy. Robyn Carston (2002) calls this strategy *The least-effort strategy (LES)*. Specifically, the hearer is doing the following when interpreting an utterance:

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 2002, p. 45)

(LES) is justified by the *Communicative Principle of Relevance*, because according to this principle, the utterance can be expected to be optimally relevant¹⁰:

The hearer is justified in following a path of least effort because the speaker is expected (within the limits of her abilities and preferences) to make her utterance as relevant as possible, and hence as easy as possible to understand (since relevance and processing effort vary inversely). ... The hearer is also justified in stopping at the first interpretation that satisfies his expectations of relevance because, if the speaker has succeeded in producing an utterance that satisfies the presumption of relevance it conveys, there should never be more than one such interpretation. (Sperber/Wilson 2002, p. 17)

According to (LES) a hearer is doing the following when interpreting an utterance: He is trying out different interpretations of the utterance, starting with the easiest ones to access. In doing this he is disambiguating, assigning referents and enriching the utterance in various ways. When he meets an interpretation that gives cognitive effects large enough to satisfy his expectation of relevance, he stops. He is justified in stopping at this point, because the speaker should formulate her utterance in a way so that the first interpretation to satisfy the hearer's expectation of relevance is the one she intended to convey (Sperber/Wilson 2002, p. 17). His interpretation may be false, since the comprehension is a non-demonstrative inference process, but “it is the best a rational hearer can do” (Wilson/Sperber 2004, p. 614).

¹⁰ See section 2.3.2

2.3.5 Identifying the implicatures

The explanation for how a hearer ends up with the implicatures he does is the same as with explicatures. By following the Least Effort Strategy the hearer will end up with an overall interpretation that meets his expectation of relevance. This interpretation will be a set of explicatures and implicatures. In the example mentioned in section 2.3.3, Peter will identify both (5) and (6) as part of what Mary conveys when uttering (4).

Implicatures fall into one of two categories: *implicated premises* and *implicated conclusions*. The implicated premises can be said to be the intended contextual assumptions, while the implicated conclusions can be said to be the intended contextual implications (Wilson/Sperber 2004, p. 615). What makes the hearer identify the implicated premises is that they “lead to an interpretation consistent with the principle of relevance, and that they are manifestly the most easily accessible premises to do so” (Sperber/Wilson 1995, p. 195). What makes the hearer identify the implicated conclusions is that “the speaker must have expected the hearer to derive them, or some of them, given that she intended her utterance to be manifestly relevant to the hearer” (Sperber/Wilson 1995, p. 195).

Returning to the example with Mary and Peter, (6) can be seen as an implicated conclusion. A possible implicated premise might be:

(10) Peter wants his food to be warm when he eats it

It should then be clear how the three interpretations (5), (6) and (10) together make Mary's utterance relevant to Peter: If Peter wants his food to be warm when he eats it, knowing that the food will be cold very soon is a reason for Peter to come and eat dinner at once. Mary's utterance will then be relevant to Peter by causing a positive cognitive effect in him.

Sperber & Wilson point out that the three sub-tasks involved in comprehension: identifying the explicatures, identifying the implicated premises and identifying the implicated conclusions. These sub-tasks should not be thought of as sequentially ordered. They are “developed in parallel against a background of expectations ... which may be revised or

elaborated as the utterance unfolds” (Wilson/Sperber 2004, p. 615).

One could say a lot more about how exactly the hearer reaches an intended meaning, and this presentation is, by far, a complete one. It should, however, be sufficient to understand C&L's arguments that will be presented in the next section. Later in this thesis, I will further expand on some of the points given in this section, and provide some additional details to RT in order to defend the theory against C&L's arguments.

3 A presentation of Cappelen and Lepore's arguments

In this section I will briefly present C&L's argument that RT cannot account for what they call 'content sharing. I will try to present it in the way it was originally presented by C&L in their two papers: "Relevance Theory and Shared Content" (2007) and "Shared Content" (2008). The reader should know that there are several points in their argument that are, at best, vague. The most important one is their notion of 'content'. I will try to interpret the more vague parts of their argument in section 4 and 5 of this thesis.

3.1 *The Non-shared content principle and Relevance Theory.*

According to C&L, Relevance Theory commits itself to the denial of 'content sharing'¹¹. They further claim that some of the central premises in RT necessarily make it so. In this part I will take a look at exactly what parts of RT C&L refer to.

In "Relevance Theory and Shared Content" (2007) Cappelen & Lepore introduce the 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.
(Cappelen/Lepore 2007, p. 3)

They further claim that RT is committed to this principle. They add that for the relevance theorists¹² the R-relation is similarity. In other words: when someone utters a sentence intending to express a certain proposition, the hearer will not grasp this proposition, but a similar one (or similar ones)¹³. I will now continue with a presentation of C&L's reasons for

¹¹ What exactly they might mean by this will be discussed in section 5.

¹² Represented by Sperber & Wilson, Carston and Bezuidenhout in their text.

¹³ Notice that in the first part of the definition of (NSC) C&L say that what is communicated is *a*

why RT, according to them, is committed to (NSC).

RT states, as explained in section 2, that when someone attends to an utterance, the first thing they meet after the decoding of the audio/visual stimuli is a *logical form* (LF) of the utterance. The LF is not propositional, and therefore not sufficient for the hearer to fully understand what has been said. To do this the hearer has to develop LF into a propositional form. Since the utterance brings its own *Presumption of optimal relevance* the hearer can use *the least effort strategy* (LES) to find the right development.

(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 2002, p. 45)

An interpretation is relevant to the extent that it has large contextual (cognitive) effects and the effort required to process it is small. Having contextual effect is to either give the audience new assumptions or strengthen or weaken old assumptions.

The problem, according to C&L, is that it is impossible to predict what proposition the hearer will end up with using this process. Since cognitive effects essentially depend on what beliefs the hearer has at a given time, and that these beliefs will vary wildly from person to person and from context to context, there is no telling what cognitive effects the utterance will produce in a given individual in a given context. There is an infinite number of possible developments of a single utterance and “common sense dictates that [different people] will all end up in different places” (Cappelen/Lepore 2007, p. 18).

The most radical illustration of this, according to C&L, is the notion of ad hoc concept construction. If a concept such as HAPPY can stand for an indefinite number of concepts and that whatever concept the interpreter latches on to depends on the interpreters cognitive states at the time of hearing the utterance, it is impossible to predict what concept he will actually end up with. The chance of it being the same as what the utterer was thinking about

proposition, while in the last part they seem to talk about the possibility of *a set* of propositions. For now, it is not important, but it will be corrected when I present a revised (NSC) in section 4.3.

seems slim (Cappelen/Lepore 2007, p. 20). What we have here is what is stated in the first part of (NSC) : The hearer has not grasped the proposition that the speaker wanted to convey.

If the hearer didn't grasp the proposition the speaker wanted to convey, what *did* he grasp? C&L claim that RT is a version of what they call “the Similarity View (SV)”. C&L elaborate this view in the following way:

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 they require is content *similarity* across contexts. (Cappelen/Lepore 2008, p. 1034)

They elaborate this further by saying that in SV the utterance 'A said that p' means the same thing as 'A said something similar to p' (Cappelen/Lepore 2008, p. 1034). This adds up to the last part of (NSC) : the hearer has only grasped a proposition R-related (similar) to the proposition the speaker meant to convey.

C&L base their claim that RT commits itself to SV on (only) two quotes from leading RT proponents. The first is from *Relevance*:

[C]ommunication 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 1995, pp. 192-3)

C&L suggest that Sperber & Wilson are saying that people don't grasp the same proposition when they communicate. They only grasp propositions that are similar to each other.

The second quote is from Carston (2001), where she explains how the *least effort strategy*(LES) helps the interpreter reach the right proposition:

[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, p. 7)

Cappelen/Lepore paraphrase this in the following way:

A speaker utters a sentence S intending to communicate the proposition that q; the interpreter 'typically' ends up with a range of propositions p1...pn, none of which is identical to q. But that's no problem, says Carston, because as long as p1...pn are similar to q, then that is good enough. (Cappelen/Lepore 2007, p. 2)

Carston's "replication in the hearer's mind" is another way of saying what Sperber & Wilson say with their expression "duplication of thought", and we see that C&L interpret it in the same way: Carston seems to claim that the speaker and hearer only grasp similar propositions. I do not think that this interpretation is correct. I will discuss this in section 4.4.

Finally, C&L quote Anne Bezuidenhout (1997)¹⁴:

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 1997, pp. 212-13)

Again it seems pretty clear that this quote indicates a SV. On the basis of this evidence, C&L claim that RT is committed to the view that "utterances of the form 'A said that p' are true just in case A said something similar to p." (Cappelen/Lepore 2007, p. 8).

These arguments may, as C&L think, imply that RT is committed to (NSC). However, one might ask: "Why does it really matter?" Why is it so important that when people communicate, the hearer will grasp the exact same proposition as the speaker wants to

¹⁴ Bezuidenhout might not, strictly speaking, be a Relevance Theorist, and so this quote might not be used as a support for their claim about a Relevance Theoretical view. However, I do think that a Relevance Theorist could agree with what Bezuidenhout is saying here, and I also think Bezuidenhout's view could be defended in the same way as I will defend the Relevance Theoretical view in a later section.

express? Why is it not enough that the hearer understands the speaker to be saying something similar to what she actually said? In the next section I will present C&L's arguments for why a theory committed to (NSC) is, as they say, “not good enough” (Cappelen/Lepore 2007, p. 4).

3.2 *The consequences*

One might think that a commitment to (NSC) does not lead to any serious consequences. C&L deny this and give several examples of what is in jeopardy if people do not share content when they communicate. In this section I will give an overview of what is at stake, according to Cappelen & Lepore (2007).

In many situations, groups of people need to **coordinate action** between them. This is often done by giving the same message to a lot of people and expect each person to follow it in the same way. For this to work we have to assume that an utterance expresses the same content in a lot of different contexts and to a lot of different people with varying cognitive states and beliefs. If people only follow a “similar” order, they will not do what is expected of them.

Not only do we care about coordinated action, but also **collective deliberation**, groups of people investigating whether something is true or not. This is something that can occur over a period of time and in varying contexts. It seems to be imperative that the question one is searching an answer to stays invariant during the whole process. C&L give an example where there is a CIA task force, consisting of several agents, trying to decide whether Igor knows that Jane is a spy. If the question about whether Igor knows that Jane is a spy means different things in various contexts and to different people, collective deliberation about this would make no sense.

The same is also true even if there is just one individual conducting an **intra-personal**

deliberation. If one individual is deliberating the same question over a period of time, it is necessary for the question to have a stable content.

A lot of our knowledge is derived from what other people say. It will sometimes form the basis for our **justified beliefs**. When we hear someone we trust say that p, it gives us good reason to believe that p. However, this only works if we are able to grasp what he said. If the person utters a proposition p and we end up believing p1, then our belief is not justified.

Another problem, connected to content sharing, is how we can hold people responsible for what they say if we do not share content. **Responsibility**, then, seems to rely on shared content. We can only hold someone responsible for what they say if we, in another context, can *understand* what they said, *say* what they said and *investigate* what they said.

Similarly: What people say often provide **reasons for action**, but it will only be a reason if one can understand it correctly. If what someone said in a different context can provide a reason for us to do something, we have to be able to *understand* this reason.

Underlying these problems are three types of data that seem to be well supported:

1. We can report what other people say
2. We can make belief attributions on the basis of what people assert
3. We can evaluate what other people say and believe

When a person utters a sentence, we can usually tell other people what that person said by uttering something like:

(11) Peter said that p

For this to work it is important that p is the same proposition that Peter actually expressed. If Peter didn't express p, but a similar proposition p' we are not right in asserting (11). The same is true in belief attribution. If Peter utters a sentence and ,by this, wants to express p',

then p' is the belief we should attribute to him. If we understand Peter to be uttering p , we will erroneously attribute p as Peter's belief. Finally, to evaluate Peter's claim that p' , the same thing is the case: we need to evaluate the proposition Peter wanted to express, not a different (although similar) one.

3.2.1 Appeal to similarity

C&L proceed with an answer that someone with a similarity view could give, to deal with these difficulties: If we say that utterances of the form 'A said that p ' are true just in case A said something similar to p , then we won't have any problems. In other words, (11) would be true even though Peter didn't express p , but a similar proposition p' . C&L present several reasons for why a view like this would not work¹⁵. I will not be discussing these reasons in this thesis, because, as I will argue in section 6.3, RT is not committed to SV. RT would not deal with the examples given by C&L by saying that: 'A said that p ' is true even if A said something only similar to p . Because of this there is no need to say anything about whether or not an account given by a SV would work.

C&L argue that any theory that implies (NSC) will have difficulties explaining how we can communicate with each other. Luckily, RT does not imply (NSC) and I will present some arguments for this in the next section. On the other hand, RT incorporates a view that implies something similar to the (NSC). I will therefore have to show how RT, in spite of this, *can* explain how we successfully communicate with each other. This will be discussed in section 6 of this thesis.

¹⁵ In (Cappelen/Lepore 2008, pp. 1035-8)

4 What Relevance Theory is committed to

C&L's central claim is that RT is committed to (NSC) and that this leads to all kinds of trouble for the theory. The problem is that (NSC) is formulated very imprecisely, which makes it too easy for RT proponents to dismiss it for the wrong reasons. In the next two sections I will show in what ways (NSC) is imprecise. I will then formulate a revised (NSC) that does not have these problems, but still presents the same challenges as (NSC) is meant to present. In the final part of this section I will argue that the interpretation, given by C&L, of the quotes from relevance theorists about “duplication of thought” is not correct. These quotes do not show that RT is committed to (NSC)¹⁶.

4.1 What a speaker wants to communicate

The first imprecision in (NSC) is the claim that: “When a speaker utters a sentence S, she is intending to communicate a proposition p”. This might be read as claiming that the only thing the speaker wants to communicate is one specific proposition p. As mentioned earlier RT says that a speaker has an *Informative intention* which is “to make manifest or more manifest to the audience *a set of assumptions I*” (Sperber/Wilson 1995, p. 58 my italics). A speaker generally wants to communicate a lot of things when uttering a sentence¹⁷. The idea that speakers only want to communicate one explicit thing when they communicate is a view that Sperber and Wilson recognize as common, but then explicitly reject (Sperber/Wilson 1995, p. 55). They illustrate this with a couple of examples. The first example is about non-verbal communication, but still shows the general point:

Mary and Peter are newly arrived at the seaside. She opens the window overlooking the sea and sniffs appreciatively and ostensively. When Peter follows suit, there is no one particular good thing that comes to his attention: the air smells fresh, fresher than it did in town, it reminds him of their previous holidays, he can smell the sea, seaweed, ozone, fish; all sorts of pleasing things come to mind, and while, because her sniff was appreciative, he is reasonably safe in assuming that she must have intended him to notice at least some of them, he is unlikely to be able to pin her intentions down any further. Is there any reason to

¹⁶ Nor the revised version of (NSC).

¹⁷ This is actually very similar to C&L's view, *Speech Act Pluralism*, which states that speakers communicate a lot of propositions when they utter a sentence (Cappelen/Lepore 2005, chapter 13).

assume that her intentions were more specific? Is there a plausible answer, in the form of an explicit linguistic paraphrase, to the question, what does she mean? Could she have achieved the same communicative effect by speaking? Clearly not. (Sperber/Wilson 1995, pp. 55-56)

Speakers communicate a lot of things, and some of those things are vaguer than others. In RT there is not really a significant difference between verbal and non-verbal communication, but since (NSC) specifically deals with spoken communication, I will provide an example that shows that also in spoken communication the speaker generally intends to communicate more than just a proposition.:

(12) Peter: What do you intend to do today?

Mary: I have a terrible headache.

It is hard to tell exactly what Mary means by this, except for the explicitly expressed assumption that she has a headache, yet there is more to her utterance; “[S]he manifestly intends Peter to draw some conclusions from what she said, and not just any conclusions” (Sperber/Wilson 1995, p. 57). It might be that she will not do anything, or do as little as possible or other similar assumptions.

It should be clear that in RT there is no assumption that what the speaker wants to communicate is just a proposition *p*. However, C&L say that there seems at least to be one explicitly communicated proposition that is derived from the words in the sentence and it is this proposition they refer to in the first part of (NSC). This interpretation of (NSC) seems plausible as in a lot of communication, like in assertions, what is communicated is mainly the proposition expressed. Another reason for thinking that C&L refer to the proposition expressed is that their argument that RT commits itself to (NSC) is based on their view of how RT explains the development of the LF. As shown in section 2, when the hearer develops the LF, the result is the hearers interpretation of the proposition expressed. I will assume that it is the proposition expressed C&L refer to when I propose the revised (NSC) in section 4.3.

4.2 What the hearer thinks he has grasped

The second imprecision lies in the last part of (NSC), namely the claim that the hearer “will interpret the speaker to have intended to communicate some proposition ... [similar] to p” (Cappelen/Lepore 2007, p. 3). This formulation is imprecise because it can be read in two different ways. Read in the wrong way one might think that it says something like: The hearer thinks that “The speaker said something similar to p” when someone utters p. This reading is done by Wedgwood in his article (2007). According to Wedgwood, C&L imply that “*similarity* is an active component of RT” and that C&L “confuse ... what interlocutors do, with the possible *effect of applying* RT” (Wedgwood 2007, p. 10-11 his emphasis). This reading is not a correct description of what C&L mean with (NSC).

According to C&L, SV implies that sentences like 'A said that p' means the same as 'A said something similar to p' and 'A understands what B said' means (something like) 'A grasped a proposition similar to the one expressed by B' (Cappelen/Lepore 2007, p. 11). However, this does not necessarily mean that the speaker himself actually thinks he has only understood something similar when interpreting an utterance. Since a view like this would be completely at odds with our own experiences with communication, it would be a strange view to attribute to any pragmatic theory. This leads me to think that C&L had a different interpretation of (NSC) in mind.

I think that a better way to understand (NSC) is to see that it is simply saying that the proposition the hearer grasps is different from (or only similar to) the proposition the speaker meant to express. It is not implying that the hearer is aware of this. It is saying something like: The hearer thinks that “The speaker said p” when someone utters p, and p' is only similar to p. This reading of (NSC) is what is suggested by how C&L define SV. I will therefore keep this in mind when formulating the revised (NSC).

It might be the case that my interpretation of C&L is wrong, and that they really *did* intend (NSC) to mean what I presented as a prima facie reading of it. If this is the case, it would

just mean that their argument could easily be rejected by RT proponents for the reasons pointed out in the last two sections. I will therefore pay no further attention to that, first, reading of (NSC).

4.3 The revised non-shared content principle

Based on the observations made in the last two sections I want to formulate a revised principle. It is just a reformulation of (NSC) as I think it was originally intended by C&L, but it gets rid of the impreciseness in their original formulation.

Revised Non-Shared Content Principle (RNSC):

When a speaker utters a sentence, S, thereby intending to express a certain proposition p, the audience will not grasp this proposition but a similar proposition p'.

Claiming that RT is committed to (RNSC) forms a more substantial challenge, as it cannot be brushed off by pointing to the imprecisions in (NSC). It also seems like the consequences listed in section 3.2 would follow as clearly from (RNSC) as from (NSC). Since (RNSC) seems to do the same job as (NSC) and it probably is what C&L meant in the first place, I will treat C&L's arguments as pertaining to (RNSC) in the rest of the thesis.

4.4 Duplication of thoughts

In this section I will take a closer look at what seems to be the reason for why C&L think that RT is overtly committed to (RNSC). By presenting some quotes from Relevance Theorists, C&L point out that in RT communication does not involve a duplication of thought. C&L think that this claim implies a commitment to the second part of (RNSC). I will argue that this is not necessarily the case.

As described in section 2.3.3 recovering the propositional form is one of the first things a hearer will have to do when understanding an utterance. Not only does he have to identify a

propositional form, but he has to identify the *right* propositional form, and that is the one meant by the speaker (Sperber/Wilson 1995, p. 183). The hearer reaches the propositional form by developing the logical form. In doing this he is doing disambiguation, reference assignment and enrichment. The whole process is guided by the principle of relevance, making the propositional form reached by the hearer highly context sensitive.

The question is, if the hearer has recovered the right propositional form, has he then not duplicated at least one of the speaker's thoughts? Namely the thought corresponding to the propositional form? The answer is "not necessarily", and the reason for this is that, according to RT, thoughts are more fine grained than propositions.

Sperber & Wilson bring up this point in their discussion of Katz' *principle of effability* which says: "Each proposition (thought) is expressible by some sentence in every natural language." (Katz 1981, p. 226). In short, Katz argues that every thought can be encoded by a unique sentence in the language, and when we meet incomplete sentences it's just because it is a convenient way of speaking. An example is the utterance:

(13) Thank God, he is gone

Katz would argue that that (13) is just a convenient way of saying something complete, like:

(14) Thank God, the man x who at time t was in location l has, at time t', left the room which the man x was in at time t.

But Sperber & Wilson say that it seems plausible that one can entertain the thought that is conveyed by (13) without entertaining anything more complete, like (14). This is because:

It seems plausible that in our internal language we often fix time and space references not in terms of universal co-ordinates, but in terms of a private logbook and an ego-centered map; furthermore, most kinds of reference – to people or events for instance – can be fixed in terms of these private time and space co-ordinates. Thoughts which contain such private references could not be *encoded* in natural languages but could only be incompletely represented (Sperber/Wilson 1995, p. 192 their emphasis)

What does this mean for the relationship between duplication of thoughts and the sharing of

propositions? Sperber & Wilson are very clear on this:

two people may be able to think *of* the same man *that* he has gone, without being able to think exactly the same thought, because they might not individuate him in exactly the same way. Similarly, by saying 'He has gone' I may induce in you a thought which is similar to mine in that it predicates the same thing (that he is gone) of the same individual, but which differs from mine in the way you fix the reference of 'He'(Sperber/Wilson 1995, pp. 192-3 their emphasis)

What follows this quote is the quote mentioned in 3.1 that C&L use to show that RT is committed to (NSC). It seems clear that RT says that one *can* grasp the same proposition without duplicating thoughts¹⁸.

One could discuss whether or not Sperber & Wilson's argument about the fine grainedness of thoughts is correct, but I do not want to do that here. The point is to show what RT overtly commit itself to, not whether or not they are correct in doing so. All in all it should be clear that one cannot use the quote from Sperber & Wilson to argue that they commit themselves to (RNSC).

However, C&L may not necessarily care about what RT *overtly* commits itself to. They argue that RT *has* to be committed to (RNSC), because of some central tenets in the theory. In section 5 I will investigate this claim. Also, even though RT is not overtly committed to (RNSC) as a general rule, it *is* nevertheless committed to it in at least some cases. I will investigate this issue in section 6 of this thesis, and argue that the consequences of (RNSC) are not as bad as C&L imagine.

¹⁸ Carston (2002, pp. 30-31) agrees with this view on thoughts as more fine grained than propositions. This means that the same argument can be made against C&L's use of the quote from Carston (2001, p. 7) as a way of proving that she is committed to (RNSC).

5 What Relevance Theory should be committed to

We have seen that RT is not overtly committed to (RNSC). C&L, however, argue that RT has some central tenets that make it impossible to deny something like (RNSC). There are certain assumptions in RT that *entail* (RNSC). In this section I will investigate what these assumptions might be. I will also expose an unarticulated argument found in C&L's argumentation. It is this unarticulated argument that makes C&L think that RT has to be committed to (RNSC).

As mentioned in section 3.1, C&L claim that since, according to RT, understanding is heavily reliant on contextual information it would be a miracle if the hearer grasped the proposition the speaker wanted to convey. What contextual information is relevant will depend on the beliefs of the interpreter, and these vary wildly from situation to situation. Because of this, C&L claim that there is no way to predict what the hearer ends up with. As Wedgwood (2007) points out, C&L seem to have missed that explaining how hearers end up with the interpretations they do is one of the main goals of RT, it is its “very *raison d'être*” (Wedgwood 2007, p. 16). I will try to make this clearer in the following section.

5.1 *Predicting what hearers end up with*

C&L are convinced that the process of bringing the hearer from the logical form to the speaker's meaning is completely random. This view does not take all the various constraints in the process of generating an interpretation seriously. I will present each of these constraints, and argue that the claim that the interpreter will end up with a totally different proposition than the speaker expressed, is too strong.

5.1.1 Constraints given by the Logical form

The first kind of constraint is the one given by the logical form. According to RT, the

logical form of an utterance contains various *concepts*. A concept consists of a label, or an address in memory, which stores different kinds of information, but it can also appear in a logical form and thereby point the hearer to the address in memory where the conceptual information is stored. The information stored in memory at a certain conceptual address fall into three types of entries: logical, encyclopedic and lexical (Sperber/Wilson 1995, p. 86). A logical entry contains deductive rules for how one can generate new assumptions from assumptions that contain the given concept. Encyclopedic entries contain information about what kind of objects the concept denotes, and assumptions about these objects. The encyclopedic entry for CAT will be a set of assumptions about cats. This entry corresponds roughly to what we would regard intuitively as the 'meaning' of a word. Finally, a lexical entry contains information about the natural-language lexical item used to express it. That is, syntactic and phonological information.

When we start developing the LF, we access the conceptual information in memory that is pointed to by the LF. So if someone utters the sentence:

(15) The cat is on the mat

my encyclopedic entry about cats will limit the possible interpretations of that sentence. Had the speaker said 'dog' instead of 'cat', I would have accessed different conceptual information. By pointing the hearer to a certain set of information, the logical form will constrain the final interpretation of the utterance,

However, the LF will not determine the interpretation completely. Usually, when someone uses the word 'cat' she means a common household cat, but in some cases the word might be used to denote related animals, like tigers or pumas. This means that (15) could be used to say that:

(16) The puma is on the mat

Furthermore, words can sometimes have a meaning that is almost completely different from the meaning it normally has, for example when they are used as metaphors. However, in all

these cases, even in metaphorical uses of words, the logical form will provide some kind of guidance for how to interpret the sentence. If someone utters:

(17) This surgeon is a butcher

what helps us understand the metaphorical use of 'butcher' is our encyclopedic entry about BUTCHERS; for example the idea that they don't cut their meat in a very precise way. This information might then be used to imply that the surgeon is careless (Sperber/Wilson 2008, p. 192).

Another example of how the logical form might constrain a hearer's interpretation, is when the utterance encodes *procedural meanings*. Utterances can, as stated above, encode conceptual representations, but they can also encode procedures for how to manipulate these conceptual representations (Wilson/Sperber 1993, p. 10). The word 'cat' would be an example of a word that encodes a concept, namely the concept CAT. The word 'so' in (18a) encodes a procedural meaning:

(18a) It is raining, so the grass is wet

(18a) communicates two propositions:

(18b) It is raining

(18c) the grass is wet

The procedural information encoded by 'so' is "Process (18c) as a conclusion" (Wilson/Sperber 1993, p. 15). Examples of other words that encode procedural meanings are: 'but', 'moreover', 'therefore', 'after all', 'you see', 'also' (Carston 2002, p. 160). Procedural meanings will in this way guide the hearer in reaching the correct proposition.

Even though the logical form puts some kind of constraints on what interpretation the hearer will end up with, it will not fully determine it. Since the concepts in the logical form may be expanded in various ways, we need other types of constraints to be able to make the right inferences.

5.1.2 Non-linguistic constraints

Related to the idea of procedural meaning is a second kind of constraint: Speakers can use non-linguistic ostensive communication along with linguistic communication, for example by gesturing or by intonating certain words. Carston argues that if it can be said that these communicative acts 'encode' anything, it would be procedural information (Carston 2002, p. 163). A nod could, for example, mean that the utterance is in agreement with something. Intonation on a certain word could mean that the hearer should pay specially attention to it, etc. In this way, the non-linguistic communication can give some additional clues that help the hearer in reaching the proposition expressed.

5.1.3 Mutually manifest facts

A third type of constraint is the fact that in a given context certain assumptions are mutually manifest. It is the speaker's responsibility to make sure that the hearer will correctly understand what she is trying to communicate (Sperber/Wilson 1995, p 43). The speaker can do this by making correct assumptions about the hearer's cognitive environment, and thereby use this information to generate an utterance that the hearer will be able to interpret correctly. The speaker can never *know* whether the hearer has a certain assumption or not, but when certain facts are mutually manifest the speaker will know that they are at least *accessible* to the hearer. The same is true for the hearer vis-à-vis the speaker:

“From assumptions about what is manifest to other people, and in particular about what is strongly manifest to them, we are in a position to derive further, though necessarily weaker, assumptions about what assumptions they are actually making. From assumptions about what is mutually manifest to all of us, we are in a position to derive further, and weaker, assumptions about the assumptions they attribute to us.” (Sperber/Wilson 1995, p. 45)

For example: Peter and Karen are in an art museum looking at different paintings, Peter then utters:

(19) I like the blue one

Peter does not need to make any guesses at what assumptions Karen might be having. He only needs to know that there is a blue painting present, and that Karen will be able to recognize this if necessary. Karen might not even have discovered it yet. She will be guided

by what is mutually manifest when she interprets the sentence. Even if Karen is thinking about which color of M&M's she prefers, and this thought is strongly accessible to her, she can be fairly certain that Peter is not referring to this private thought when uttering (19), but to the mutually manifest blue painting.

Just as with the constraints given by the logical form, this constraint cannot *guarantee* successful interpretation. Considering that none of the participants *know* what assumptions the other person is actually making, there is room for mistakes. Let us say that in the example above, Karen has been thinking so hard about M&M's that she has failed to notice the blue painting. She may have been so focused on her own thoughts about M&M's that she believes that Peter is also talking about M&M's when uttering (19).

5.1.4 Contextual constraints

The fourth type of constraint is the choice of context. Even though the context might theoretically include all of the hearer's assumptions, this is rarely the case. The hearer's context will generally consist of only some of his assumptions. This further constrains the possible interpretations of a sentence.

Sperber & Wilson say that, contrary to what one might think, the context is not completely determined in advance of an utterance interpretation. The hearer first assumes relevance and then he chooses a context which will maximize relevance: “[R]elevance ... is treated as given, and context ... is treated as a variable” (Sperber/Wilson 1995, p. 142). There are of course some assumptions that can be treated as given. These assumptions belong to the *immediately given context*. It consists of the set of assumptions used in interpreting the last utterance and other assumptions in the hearers short-term memory store. This immediately given context can, however, be extended in various ways; for example by including interpretations of utterances preceding the last one; by adding encyclopedic entries for concepts present in the utterance and adding information about the immediately observable environment (Sperber/Wilson 1995, pp. 140-1).

The point is that the context is extended *as needed*. Unless new information seems relevant, the hearer will not expand his context beyond the initially given one. This makes it much more predictable what kind of contextual effects an utterance might have. This is not only because the amount of assumptions included in the hearer's context is limited, but also because a lot of the assumptions are mutually manifest and thus, shared by the speaker.

5.1.5 Constraint given by the Communicative Principle of Relevance

The last type of constraint I want to discuss is what is given by the Communicative Principle of Relevance: every utterance conveys a presumption of its own optimal relevance. C&L are aware of this principle, but they fail to see how it helps the hearer in recovering the right meaning.

Remember that, according to RT, when someone wants to communicate something, what she wants to convey is a set of assumptions **I**. The Communicative Principle of Relevance provides the hearer with one of the assumptions in **I**, namely the presumption of relevance. This is helpful because “[t]he presumption of relevance is not just a member of **I**, it is also *about I*. As a result, it can be confirmed or disconfirmed by the contents of **I**.” (Sperber/Wilson 1995, p. 164).

The presumption of relevance says two important things about the assumptions the speaker wants to communicate: The assumptions are relevant enough for the hearer to be justified in using some effort in processing the utterance, and the utterance is the most relevant one the speaker could use to communicate those assumptions (for the complete definition, see section 2.3.2).

Any assumption the hearer comes up with in the processing of the utterance, will be checked against the information given by the presumption of relevance. If a certain assumption is either not relevant enough, or it could have been communicated better with

another utterance, the hearer is justified in thinking that it is not an assumption the speaker wanted to convey. Because the propositional form is what is supposed to guide the hearer to the right assumptions, and the presumption of relevance constrains what assumptions the hearer ends up with, the propositional form itself is constricted by the presumption of relevance.

5.1.6 The constraints cannot guarantee success

I have shown that there are a lot of mechanisms that both constrain the range of possible interpretations of an utterance, and also guide the hearer to make the right inferences when he is interpreting the utterance. Given the above, the statement from C&L that “[w]e have *no way* to predict in advance which development of these logical forms [the hearers] will end up with” (Cappelen/Lepore 2007, p. 18; my emphasis), seems to be false.

Having said all this, it is important to keep in mind that the constraints still do not guarantee that the hearer will end up with the intended assumptions. Even after the constraints have done their job, there might not be any *one* given proposition expressed. This is especially clear in metaphorical uses of language. Depending on the context, even after applying all the constrictions, a hearer might end up with various interpretations of (17):

(17) This surgeon is a butcher

What *exactly* is meant by 'butcher' may not be totally obvious, and the speaker and hearer might not have the exact same meaning of the term 'butcher' in mind. This point might also apply to more normal uses of language. I will be discussing this in section 6, where I argue that one can communicate successfully even when the speaker and hearer do not interpret an utterance in the same way.

5.2 The unarticulated argument

C&L claim that there is no way to predict what the hearer will end up with. I think the reason for their claim can be found in an unarticulated argument. In this section I will

uncover that argument, and show that if it is sound, then C&L are right in claiming that RT has to be committed to (RNSC).

To expose the hidden argument, it is helpful to look at something C&L say in their paper (2007). C&L claim that RT is not the only theory that endorses (NSC). (NSC) is actually a view that is endorsed by *any* theory that adheres to three principles, (RC19)-(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. (Cappelen/Lepore 2007, pp. 4-5)

Based on RC1-RC3 it seems that C&L require a sentence to semantically express a proposition if one wants to avoid having a view that implies (RNSC). If there is no invariant proposition, we cannot understand what other people say, only something similar.

Why do C&L assume this? It seems that C&L believe that there has to be a *guarantee* in communication that the hearer ends up with the same proposition as the speaker¹⁹, and that without such a guarantee the hearer will never be able to recover the speaker's meaning. This is clear from C&L's comments about the *Least Effort Strategy*(LES): “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)” (Cappelen/Lepore 2007, p. 19). It is also evident in the following claim:

The proposition semantically expressed is our minimal defense against confusion, misunderstanding, mistakes and it is that which *guarantees* communication across contexts of utterance. (Cappelen/Lepore 2005, p. 185 my emphasis)

The only way to provide a guarantee for successful communication would be, according to C&L, that the proposition expressed is determined by the semantics of the sentence.

¹⁹ What such a guarantee amounts to will be discussed in section 5.4.1.

Based on the above observations, I will formulate what seems to be C&L's unarticulated argument:

- P1: The proposition expressed must be determined by the semantics of the sentence for the hearer to have any chance in recovering it.
- P2: RT states that the proposition expressed is pragmatically determined.

Conclusion: Therefore, given RT the hearer will never recover the proposition expressed.²⁰

Since the argument is valid, and since P2 is true, I will have to show that P1 is false in order to deny the conclusion. I will try to do that in the remainder of section 5.

5.3 Content

C&L's notion of 'semantically expressed' is, at best, misleading. Initially, one may think that C&L mean something that is determined completely independent of context. I will show that this is not the case. Understanding what they mean by 'semantically expressed' relies on understanding what they mean by 'content', since these two notions are related in their discussion. Because of that, I will start by investigating what they mean when they use 'content'. After that I will reformulate P1 to avoid the impreciseness pertaining to C&L's use of 'semantically expressed'.

5.3.1 Cappelen & Lepore's use of 'content'

The term 'content' is used in different ways in the literature²¹, so it is important to be specific about what definition one is committed to when using it. For C&L this is especially important because their claim is that RT cannot account for 'content sharing' across context.

²⁰ Because of this, RT forced to accept (RNSC).

²¹ As pointed out by Korta & Perry (2006: pp. 451-2).

Unfortunately C&L are not as clear as they should be, leading to terminological confusion in the debate. My first task will, therefore, be to clear up the terminological confusion in order to work on the real issue.

A good start would be to look at (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 2007, p. 3)

Since the principle is supposed be about how people do not sharing content, it might be helpful to understand exactly what it says. However, (NSC) as formulated by C&L, is not clear about what the term 'content' might refer to, because it does not appear in the definition. As argued in section 4 there are other difficulties with the formulation of this principle, and I argued that what C&L mean by (NSC) is better expressed in what I called (RNSC):

Revised Non-Shared Content Principle (RNSC):

When a speaker utters a sentence, S, thereby intending to express a certain proposition p, the audience will not grasp this proposition but a similar proposition p'.

Neither in (RNSC) is the term 'content' used, but one might suggest that 'content' refers to the proposition expressed. The problem is that the term 'the proposition expressed' has its own difficulties, as will be shown in section 5.3.2. I will therefore try to provide some other clues to what C&L might mean with their use of 'content'. However, one might at least be justified in thinking that what they refer to is propositional.

In *Insensitive Semantics* (2005) C&L define what they call *semantic content*: “The semantic content of a sentence S is the content that all utterances of S share.” (Cappelen/Lepore 2005, p. 143). They further add that it is the content that can be grasped by someone not familiar with the context of the utterance. In other words, it seems like the semantic content is what you grasp if you know what the words in the sentence mean, and know the grammar

of the language the sentence is expressed in. This content can be grasped without any knowledge of the context. In RT, this level of content is called the *logical form*, or the *encoded meaning* of the utterance. However, this level of content does not appear to be propositional. If the sentence contains ambiguous terms or terms that need reference assignment, the truth conditions of that sentence will vary between contexts.

In their article “Shared Content” (2008), C&L use the expression 'what is said' to refer to the content of an utterance. This is not helpful as it is unclear what exactly 'what is said' is referring to. In our everyday use it might be used to refer to anything from the exact words used in the utterance (E.G. when one quote someone), to the implications of the utterance.²²

As I have tried to show, C&L's use of 'content' has been somewhat imprecise. This has led to several misunderstandings in the discussion on shared content²³. Before I continue with my discussion, I will define several types of content in order to avoid the terminological confusion that is present in the debate. This will make it easier to pinpoint how C&L use the term 'content' and also see how their use differs from other uses of the term.

5.3.2 Defining content

In this section I will define four notions of content. They each point to a different level of content in an utterance. They also overlap, so the fourth level of content contains the three previous ones, and the first level is contained in all the other levels.

The first kind of content, I will refer to as '**semantic content**'. A definition for this level can be found in C&L's book (2005): “The semantic content of a sentence S is the content that all utterances of S share.” (Cappelen/Lepore 2005, p. 143). This level is what is known as the logical form, or encoded meaning, in RT. By understanding what the words in the

²² For a more substantial discussion on the term 'what is said' see Recanati (2001)

²³ E.G. in Korta/Perry (2006) and Wedgwood (2007). This misunderstanding is also pointed out by C&L (2006) and Kriempardis (2009: pp. 167-8)

sentence mean and how the grammar of the language works, one knows what the semantic content is. It is important to point out that this is *not* how C&L later use the term 'semantic content'²⁴. What C&L later call 'semantic content' is something over and above what is shared by all utterances of a sentence S. This is also pointed out by Carston (2007).

The next level of content I will call '**Filbert content**', as suggested by C&L (Cappelen/Lepore 2006, p. 31). What C&L refer to when they speak about 'semantic content' (or as they some times call it, 'minimal semantic content') seems a little more context dependent than what I have defined as semantic content. There is *some* pragmatic influence on C&L's 'semantic content'. This is also pointed out in "Precis of Insensitive semantics": "Of course we realize that [their notion of semantic content] is not stable with respect to the values of indexicals and demonstratives, ambiguous and vague expressions, etc." (2006, p. 29). In their reply to Korta&Perry C&L admit that their notion of 'content' is different from what others call 'minimal', since it requires these different kinds of contextual influences (2006, p.31) .

C&L define Filbert content explicitly in their reply to Korta&Perry (2006, p. 29). C&L say that this level of content is what is characterized by points (a)-(d) on their list in *Insensitive Semantics* (2005):

- (a) Specify the meaning (or semantic value) of every expression in S ...
- (b) Specify all the relevant compositional meaning rules for English ...
- (c) Disambiguate every ambiguous/polysemous expression in S.
- (d) Precisify every vague expression in S. (Cappelen/Lepore 2005, pp. 144-5).

Points (a)-(b) is what corresponds to what I call semantic content. It should be clear that when adding (c)-(d) one gets something which is not shared by every utterance of a sentence. The ambiguous sentence:

(20) John went to the bank

is an example of this. An utterance of (20) (ignoring the need for a reference assignment of

²⁴ This is probably the major reason for the terminological confusions in the debate.

'John') will in some cases be used to say that John went to a financial institution, and in other cases say that he went to a river bank. C&L claim in their reply to Korta&Perry that the level of content characterized by the points (a)-(d), is the first level of content that is propositional (Cappelen/Lepore 2006, p. 29)²⁵.

By taking C&L's Filbert content characterized by (a)-(d) on their list and add the extra element (e): "Fix the semantic value of every context sensitive expression in S" (Cappelen/Lepore 2005, p. 145) you get what C&L call 'the proposition semantically expressed'. I will not use this term, but call this level of content '**Fixed Filbert content**', partly because there is a disagreement whether this level of content is propositional or not (see further down). Be aware that even though C&L claim that the proposition is "semantically expressed" one actually needs to do some contextual enrichment to get this level of content. This is something C&L are aware of: "The minimal proposition cannot be characterized *completely* independent of the context of utterance" (Cappelen/Lepore 2005, p. 143).

Fixed Filbert content is, according to C&L, always propositional. According to RT it is most of the time in need of further enrichment to become propositional. An example of this would be the sentence:

(21) This fruit is green

C&L claim that this is propositional (after assigning a referent to 'this'). RT claims that it is not propositional. Carston gives an example where you may use (21) on different occasions to either talk about the peel of the fruit or the interior of the fruit. Only when this information is supplied you have a proposition, according to her (Carston 2002, pp. 23-24).

²⁵ Unless there is a context sensitive expression in the utterance. For Cappelen & Lepore only a few expressions in the language are context sensitive. For example the personal pronouns and demonstratives (Cappelen/Lepore 2005, p. 144). They say nothing about proper names, as with the use of 'John' in (20), so it is unclear how they would treat this. To me it seems obviously context sensitive (since there are a lot of 'Johns'), so they would probably treat it in the same way as other context sensitive expressions.

The first level of content that is propositional, according to RT, is the explicature of the utterance. I will refer to this level of content as the **explicated content**. You get to this level after taking the fixed filbert content and enriching the sentence even further. An example of the explicated content of (21) would be:

(22) The interior of this fruit is green

I will not enter into a discussion on exactly when something becomes propositional. For my purposes it is just important to be clear about how C&L use the different terms. In this section I have shown how C&L's use of some important terms is different from other people's use of those terms. I will use this information to revise the 'unarticulated argument', first presented in section 5.2, in an attempt to avoid confusion about what exactly C&L seem to claim.

5.3.3 Revised version of the unarticulated argument

Seeing how C&L use the terms 'semantically expressed' and 'the proposition expressed', we can get a clearer view of what the unarticulated argument amounts to. A closer inspection reveals that there are several similarities between RT and C&L's view. If we take 'the proposition expressed' to refer to fixed filbert content, there seems to be no disagreement between RT and C&L's view. Neither C&L nor RT claim that the fixed filbert content is determined by semantics alone. Both positions agree that there is some pragmatic influence on the determination of this level of content. C&L and RT also agree that this level of content is shared between speakers. The only disagreement between the two positions is that C&L claim that this level of content is the first level of content that is propositional, and that RT claims that this level of content is only sub-propositional²⁶.

The reason why C&L claim that RT is committed to (RNSC) is that according to RT, 'the proposition expressed' is not determined by fixed filbert content, but needs further enriching. This further enriching is what makes it impossible for a hearer to grasp the same proposition as the speaker expressed. The first premise in the unarticulated argument can then be reformulated as:

²⁶ As mentioned: In RT, the first propositional level of content is explicated content.

P1: The proposition expressed must be determined by fixed filbert content for the hearer to have any chance in recovering it.

In section 5.4 I will investigate what reasons C&L might have for thinking that P1 is true. First I will present some arguments for how C&L can allow *some* pragmatic influence on their 'proposition expressed', while at the same time claim that RT's notion of 'proposition expressed' allows *too much* pragmatic influence.

5.3.4 The difference between fixed filbert content and explicated content

We have seen that C&L allow the possibility of *some* pragmatic influence on the proposition expressed, but at the same time disallow *too much* pragmatic influence. Why is explicated content any worse off than fixed filbert content as a candidate for 'the proposition expressed'? In this section I will try to provide an answer to this question.

The reason for C&L's acceptance for certain types of pragmatic influences can be found by noticing that there are two different kinds of pragmatic processes that contribute to explicatures. They are referred to as 'saturation' and 'enrichment'. These processes both tell the hearer to add certain contextual information to the utterance to make it complete. *Saturation* involves "finding the intended content (or value) for a linguistically indicated variable or slot" (Carston 2009, p. 49). In other words, there are elements in the sentence that indicates that the hearer will have to do some contextual 'filling' to fully understand them. A clear example of this is whenever someone uses an indexical, like the words 'he' or 'that'. Both words clearly signal that the hearer will have to find some kind of contextual information to fill out the values in these slots.

Free enrichment, on the other hand, does not involve filling out syntactically given 'slots'. The linguistic form does not give any indication that free enrichment is needed (Carston 2009, p. 49). An example, would be the sentence:

(23) He has a brain

which, in normal circumstances, usually would be taken to mean:

(24) He has a high-functioning brain

However, there is nothing in this sentence itself that suggests to the hearer that enrichment is needed. Indeed, there might be contexts where (23) does not need further enrichment (for example some kind of Frankensteinian-monster scenario, where the participants discuss what body parts they have given to the monster).

It may be helpful to think of saturation as a 'bottom-up' process, and free enrichment as a 'top-down' process. Both processes involve adding contextual information to the utterance, but saturation is linguistically mandated while free enrichment is context-driven (Recanati 2004, p 23).

Some people do not think that top-down processes like free enrichment contribute anything to the explicature of the sentence. Jason Stanley (2002) argues that the only pragmatic process needed for grasping the explicature, is saturation²⁷. Stanley thinks that processes like free enrichment are so unconstrained that nothing would hinder the hearer in constructing a meaning completely different than what the speaker had in mind. Hall (2008) gives an argument, similar to mine, about how the Relevance theoretical process have several constraints that will hinder this type of 'over-generation'. Whether or not Stanley's claim is correct, we can still see why he is reluctant to allow top-down pragmatic processes: these processes are, in his view, unconstrained, and to appeal to these processes is “to abandon hope of giving any systematic explanation of how we communicate linguistically” (Hall 2008, p. 428)

The thought that free enrichment is 'unconstrained' seems to be the reason for why also C&L allow saturation, but not free enrichment, to play a part in determining the proposition expressed. C&L's also believes, as Stanley does, that free enrichment is unconstrained. This is what they point to when they claim that it is impossible to tell how the hearer will develop the LF. This belief is also why C&L allow saturation, but not free enrichment, to

²⁷ I will not be discussing his argument in depth in this paper as it is not relevant for my argument. For an in depth discussion of Stanley's argument, see Hall (2008).

play a part in determining the proposition expressed. The bottom-up processes not only explicitly 'tell' the hearer to fill in some contextual information, but they also give specific guidelines to what the hearer must search for (E.G. filling the position marked by 'he' would most probably be some salient individual in the context)²⁸.

5.4 Why is propositional fixed filbert content necessary?

Hopefully, I have shown why one might think that fixed filbert content is easier for the hearer to grasp than the explicature. We can now return to the question I first set out to answer: Why is it so important that the proposition expressed can be determined by the semantics of the utterance alone (or the semantics plus some minimal, bottom-up, pragmatic influence). Why can we not also allow free enrichment to play a part in determining the proposition expressed? I will argue that C&L's reason for not allowing free enrichment to play a role in determining the proposition expressed, is that they require some kind of guarantee of successful communication. C&L require a level of content that is always shared between the speaker and hearer, even if their respective contexts are wildly different. I will argue that a guarantee of successful communication cannot be had, and that C&L's fixed filbert content cannot do what C&L require of it.

5.4.1 The need for a failsafe procedure

As mentioned in section 5.2 it seems like C&L want a procedure that is failsafe, a procedure that can *guarantee* success. This requirement appears strange for a reason pointed out by Sperber & Wilson (1995): “Since it is obvious that the communication process takes place at a risk, why assume that it is governed by a failsafe procedure?” (pp. 44-5). We know that errors are extremely common in communication. People misunderstand each other a lot. The data should point us in the direction of a procedure that not only explains how we understand each other, but also why we so often *misunderstand*

²⁸ One could argue that even saturation cannot completely determine what the hearer ends up with. If someone utters “I will stay here”, exactly what is meant by “here” can be hard to determine. How far can the speaker move from her current position without rendering her utterance false?

each other. RT explains how people can communicate, but also shows how fragile the process of understanding is:

[E]ven under the best of circumstances ... communication may fail ... The best [the hearer] can do is construct an assumption on the basis of the evidence provided by the communicator's ostensive behaviour (Sperber/Wilson 1995, p. 65).

A hearer can never be *completely* certain that his understanding of what the speaker meant is correct, but depending on the clues given by the communicator, he can be more or less certain. After all, if communication were guaranteed to be successful, we would never misunderstand each other. RT “provides a solution which is, of course, imperfect, but is nonetheless effective” (Wilson/Sperber 2002, p. 606). Because of this, believing that C&L require a guarantee for successful communication at *every level* is not reasonable. However, they do say that at a certain minimal level you *can* have a guarantee of success. This level is what they call the minimal proposition, or what I call propositional fixed filbert content.

Elugardo (2007) expands on this claim. He argues that C&L believe the minimal proposition takes the role of a 'cognitive safety-mechanism'. When hearers are ignorant of, or mistaken about, certain features of the context, they can always fall back on the minimal proposition. They can have a guarantee that the speaker at least meant this proposition, and likewise, the speaker can be certain that the hearer will at least grasp this proposition.

The first problem with this claim is that there is no reason to think that even at the level of the minimal proposition (fixed filbert content), can you have a guarantee of success. There is no guarantee that the people involved grasp the same content, because there is room for errors even before you reach fixed filbert content.

To be able to grasp the fixed filbert content, the hearer first needs to grasp the semantic content. The semantic content is always shared by the speaker and hearer as long as they both know the language they speak. The hearer then needs to process the semantic content by disambiguating every ambiguous/polysemous expression and precisifying every vague expression in the sentence. How to do this will be decided by the context of the utterance,

and, because of this, the hearer is at a risk of not doing it correctly. After doing this, the hearer has to fix the value of every context sensitive expression in the sentence. How to do this is also, obviously, decided by the context and, again, the hearer is at risk of not doing this correctly. Wedgwood says that “given an utterance of *He's ready*, surely a response like *Who? Mark or Paul* is at least as likely as *Ready for what?*” (Wedgwood 2007, p. 21 his emphasis). Wedgwood goes so far as to say that “[resolving context sensitive expressions] is one of the more obvious sources of misinterpretation in language” (Wedgwood 2007, p. 21).

There is another problem with C&L's reliance on the minimal proposition. Even if we were to grant that this minimal proposition had a guarantee of being successfully shared, it would still not do the work C&L claim it does, namely serve as a safety mechanism. The hearer is *not* justified in believing that the speaker at least meant the minimal proposition when they lack contextual information, as C&L claim. Nor is the speaker justified in thinking that the speaker will grasp this minimal proposition. I will provide some arguments for this in the next section.

5.4.2 C&L's minimal proposition does not function as a safety mechanism

C&L claim that their minimal proposition (fixed filbert content) serves as a safety mechanism in communication. In situations where hearers and speakers are in different contexts and the hearers are ignorant of relevant contextual facts in the speakers environment, the hearer can still safely assume that the speaker at least meant the minimal proposition of her utterance and the speaker can expect the hearer to at least grasp this minimal proposition (Cappelen/Lepore 2005, 184). Let us assume that A utters:

(25) John is ready

B, who is in a different context, is being told that A uttered (25). Now, since B does not know anything about the context in which A uttered (25) he cannot completely understand the utterance²⁹. However, he does know *one* thing: namely that John is ready, simpliciter.

²⁹ Since he does not know the answer to the question “ready for what?”

C&L claim that there is something in common in every use of the word 'ready': a state of readiness (Cappelen/Lepore 2005, p. 167). The hearer can know that John is in a state of readiness, without knowing exactly what John is ready for, and this is the minimal proposition expressed by (25). In this way John can at least start to understand what has been said by (25). The minimal proposition serves as a “starting point” for further discussion (Cappelen/Lepore 2005, p. 185).

If this is right, then one could agree that C&L's minimal proposition plays an important role in communication. However, Borg (2006) and Wedgwood (2007) give several reasons for why the minimal propositions does not play the central role that C&L want it to play.

Remember that C&L give some examples (mentioned in section 3.2) of communicational situations where they believe that shared content is indispensable. If the minimal proposition does the work C&L claim it does, they need to show that it is the sharing of *this* proposition that makes the communication successful. Wedgwood (2007), however, provides some examples showing that this is not the case. What needs to be shared is not the minimal proposition, but something that has been further enriched.

An example, given by C&L, of a situation that demands sharing of content is when we hold people responsible for what they have said. However, C&L's minimal proposition might not be able to play the part as this shared content. The following example should make this clear. Let us say that someone is building a house and is in doubt whether the steel is strong enough to support the roof. She asks the engineer about this, and the engineer answers:

(26) Steel is strong enough

According to C&L, the engineer would then only be taken to be responsible for the minimal proposition that steel is strong enough (simpliciter). This claim is clearly wrong, as seen if the roof they were talking about collapses and the engineer says: “I only said steel is strong enough. I never said strong enough to support the roof”. Clearly, in this case, we would hold the engineer responsible for more than just the minimal proposition (Wedgwood 2007,

p. 26). Another example would be if a person excuses herself from football practice by uttering:

(7) I have a temperature

Clearly we would not think the minimal proposition conveyed (which is true as long as she has *any* temperature) would be a good enough excuse to skip football. The speaker would be held responsible for saying that she has a temperature above normal. Similar arguments can be made for the other examples given by C&L³⁰.

What about the claim that the minimal proposition serves as a fall back position? Even though the engineer when uttering (26), is committed to something more than the minimal proposition expressed by it, it seems reasonable to think that the hearer can at least be sure that the engineer *at least* meant this minimal proposition. Emma Borg (2006) argues that not even this claim is something that holds in all situations. Borg gives an example of a person looking into the fridge while uttering:

(27) There is nothing to eat

Borg claims, correctly, that the speaker “would be surprised to learn that they had asserted the quite general proposition that *there is nothing to eat (in the universe)*” which is the minimal proposition expressed by (27) (Borg 2006, p. 16). The speaker would probably even deny that she had asserted this minimal proposition. One may also assume that a rational hearer would never think that the speaker was committed to the minimal proposition expressed by (27). An even better example of speakers not meaning the minimal proposition would be cases of irony. If someone say ironically:

(28) Jill was nice

The speaker does *not* want to assert the proposition that Jill was nice, and neither should the hearer think so.

Borg also note that there is experimental evidence showing that minimal propositions does

³⁰ Wedgwood (2007) gives a few more. It should be clear enough, however, how easy it is to construct an example where the minimal proposition does not do the work C&L claim it does.

not always play a part in the hearers comprehension of an utterance. For example there is evidence that a hearer might recognize the need for a metaphorical interpretation of a word without first processing the entire sentence to recover the proposition literally expressed (the minimal proposition) (Borg 2006, p. 16).

I have shown that there seems to be no reasons to think that fixed filbert content has a guarantee of being successfully shared. The only level of content that possibly has this guarantee is the semantic content (Wedgwood 2007, p. 3)³¹. Even if we, for the sake of argument, grant that there is a guarantee that the hearer will end up with the same fixed filbert content as the speaker, this would still not be of much help. Fixed filbert content does not necessarily play any important part in the comprehension process. The speaker might not even mean to assert what is conveyed by the fixed filbert content. The conclusion is that fixed filbert content cannot do what C&L require from it.

5.5 *Where do we stand?*

We have seen that C&L's fixed filbert content cannot provide a guarantee that the speaker and hearer will end up with the same proposition. There is no reason to think that *any* level of content can do this. If this is true and P1 of the unarticulated argument is true, then one *has* to be committed to RNSC; Hearers never grasp the proposition expressed by the speaker. They will at most grasp a similar proposition.

Whether or not RT is committed to RNSC then relies on the truth of P1. C&L seem to think this premise is true because they want a guarantee of success, and I have shown that such a

31 Even this might not be entirely correct. As Sperber & Wilson say: "The fact that a public word exists, and is successfully used in communication, does not make it safe to assume that it encodes the same concept for all successful users" (Sperber/Wilson 1998: p. 200). This is an interesting question that requires its own discussion, but is somewhat outside the scope of this paper. However, even if one assumes that not even the semantic content can be perfectly shared (which seems to be what RT, in fact, does), RT will still be able to explain how communication works without any change in the theoretical framework. One would just have to treat every word as an ad hoc concept.

guarantee can not be had. However, even though we cannot have a guarantee of success, it does not follow that speakers and hearers *never* grasp the same propositions. C&L fail to see that RT tries to explain how people, in many cases, *do* end up with the same propositions, even though there is no guarantee that they will. RT does not provide a procedure that is failsafe, but as shown, a failsafe procedure would not be able to account for frequent misunderstandings in communication. RT gets its strength from being able to explain why we so often understand each other despite the fact that we only have partial evidence of what the speaker means.

However, even if it does not follow that a lack of guarantee will lead to only a similar understanding in the speaker and hearer. There are examples where it seems really plausible that the speaker and hearer only grasp similar propositions. An obvious example would be in cases of ad hoc construction of word meaning. If Gudrun utters the words:

(29) Svanhild is depressed

the hearer might construct a meaning of 'depressed' to be something like: "A general feeling of sadness". Gudrun, because she is a psychologist, also meant that Svanhild's appetite is low; that she does not get enough sleep and other physical/behavioral properties attributed to the clinical definition of 'depression'. It is plausible that in this case, the hearer would construct a slightly different proposition than what Gudrun expressed.

RT happily agrees with this, but claims that this does not necessarily mean that the communication was unsuccessful. The speaker and hearer could, in many cases, continue their conversation even though they formed a slightly different understanding of (29). In the next section I will investigate the possibility of successful communication even when speakers and hearers do *not* share propositions. I will argue that, in spite of C&L's arguments, one *can* talk about the same thing, collectively deliberate ideas etc, even when the people involved only grasp similar propositions.

6 Communication without content sharing

Up to this point I have been arguing that speakers and hearers *do* share propositions, but that the propositions are determined in context and not by semantics alone. In this section I will investigate the possibility of communication in situations where people do not share propositions. I will present RT's view that successful communication does not rely on grasping the exact same propositions, and that in many cases one can get by with only a similar proposition. I will then argue that this does not amount to what C&L call 'the Similarity View'. In the last part of this section I will examine C&L's examples of communication (given in section 3.2) where, according to them, sharing of propositions is necessary for communication to be successful. I will argue that this is not the case, and that even in the examples presented by C&L, communication can be successful without speakers and hearers sharing propositions.

6.1 Resemblance

What we have is a dilemma: It seems that in many cases, people do *not* grasp the same proposition when they talk with each other. An example of this is when Gudrun uttered (29) and the hearer constructed a slightly different meaning of 'depressed' than what Gudrun had in mind. On the other hand, it seems like that in the same cases, the people involved actually *are* talking about the same thing (to some extent). Depending on what Gudrun and her audience do with the information in (29), the slight difference in the meaning of (29) may never matter.

The key to explaining this is to see how, according to RT, two propositions can *resemble* each other: "Two propositions resemble each other in a given context to the extent that they share logical and contextual implications in that context" (Wilson 1995, p. 208).

'Contextual implications' are defined in "Loose talk" (1985/1986):

In a context {C}, a proposition P may have what we call *contextual implications*. A

contextual implication of P in the context {C} is a proposition implied neither by {C} alone, nor by P alone, but by the union of {C} and P. (Sperber/Wilson 1985/1986, p. 157)

This idea of *resemblance* is expanded in Wilson/Sperber (2002). They describe a situation where Mary and Peter are deciding on where to go on their next bicycle trip, and Peter has just said that he feels rather unfit. Mary utters the sentence:

(30) Holland is flat

wanting to convey that Holland might be a good choice, given Peter's physical condition. The first thing to notice is that in (30) the word 'flat' is used *loosely*. Mary does not mean to indicate that Holland is literally flat. She is indicating a concept FLAT* which meaning “is such that Holland's being FLAT* is relevant-as-expected in the context” (Wilson/Sperber 2002, p. 614). In this example a rough answer would be that a terrain is FLAT* if traveling across it involves little or no climbing. However, this is not supposed to say exactly what Mary means. What she actually means is something more vague:

The claim that Holland is FLAT* carries a range of implications which Mary expects to satisfy Peter's expectations of relevance. The concept FLAT* is individuated (though not, of course, defined) by the fact that, in the situation described, it is the first concept to occur to Peter which determines these implications. (Wilson/Sperber 2002, p. 616)

Mary might not even herself have a clear idea of the exact meaning of FLAT*³²:

The implications which Mary expects Peter to derive need not be individually represented and jointly listed in her mind. In normal circumstances, they would not be. (Wilson/Sperber 2002, p. 616).

Since the loose use of 'flat' in Mary's utterance makes the meaning of the utterance so vague, it would not be unexpected if Peter constructed a slightly different concept, FLAT**. He would then have grasped a proposition which was slightly different from the proposition Mary meant to express. However, since these two propositions resemble each other by having roughly the same import in the situation, Wilson & Sperber say that: “This would not be a case of imperfect communication or insufficient understanding” (Wilson/Sperber 2002, p. 617). This appears to be right, since Peter would correctly have understood Mary to be suggesting that they should go to Holland, since it is flat enough for Peter to ride a bicycle there, even though he is unfit. Note that this type of resemblance is

32 I will elaborate this claim in the next section.

context-dependent. If the situation were different, the two propositions might not have resembled each other closely enough (Sperber/Wilson 1985/1986, p. 157). This is because, as mentioned above, contextual implications are determined not only by the proposition, but also by the context in which it is being uttered.

Sperber & Wilson correctly claim that this 'looser' kind of understanding is what mostly, in ordinary communication, is intended and achieved. They say that:

the type of co-ordination aimed at in most verbal exchanges is best compared to the co-ordination between people taking a stroll together rather than to that between marching in step. (Sperber/Wilson 1998, p. 199)

In the next section I will take a closer look at exactly what RT claims is being “aimed at” in normal communication.

6.2 What is the aim of communication

My starting point will be *The Informative Intention* which says that a speaker intends to “make manifest or more manifest to the audience a set of assumptions **I**” (Sperber/Wilson 1995, p. 58). A way to do this would be to express a proposition that will make the hearer derive the assumptions as logical or contextual implications of that proposition, as long as the hearer can be expected to ignore the implications that are not intended by the speaker (Sperber/Wilson 1985/1986, p. 162). Sperber & Wilson give an example, showing how a literally false proposition can do this job (Sperber/Wilson 1985/1986, p. 163): Marie lives in Issy-les-Molineaux, a block away from the city limits of Paris. At a party in London a person asks her where she is from and she answers:

(31) I live in Paris

Marie's answer is, as Sperber & Wilson point out, literally false, but in the circumstances described the answer would yield all the implications Mary wanted to communicate: for example that she spends most of her time in the Paris area; that she knows Paris, that she lives an urban life; etc. Uttering (31) was the most economical way of conveying these

implications; A longer, more correct, description would require more processing effort without any more cognitive effect. Actually, by uttering something a little more true like:

(32) I live near Paris

Marie would mislead Peter into thinking of a bunch of false propositions. Peter would presumably infer that Marie would have to travel some distance to get to Paris, that she lives a suburban life etc. The surprising conclusion is that: “Marie's first answer ... is effective enough to convey just what she wants; it may be more effective than the literally true second answer” (Sperber/Wilson 1985/1986, p. 164). The moral is that the proposition expressed, might not be what matters in communication, it might just be a means to convey a bunch of other propositions. If this is right, then it does not matter if the hearer grasps the proposition expressed, as long as he grasps the propositions the speaker wanted to convey. This seemed to be the case in interpreting (30). It did not really matter if the hearer grasped the same proposition as the speaker, since the slightly different concept, FLAT**, could be used to infer the propositions the speaker wanted to convey.

Another, related, point is that what the speaker wants to communicate is, most of the time, partly precise and partly vague. As mentioned in the last section, the speaker does not necessarily have a clear idea of exactly what assumptions she wants to convey. Sperber & Wilson say:

[T]he communicator must have in mind a representation of the set of assumptions **I** which she intends to make manifest or more manifest to the audience. However, to have a representation of a set of assumptions it is not necessary to have a representation of each assumption in the set. Any individuating description will do. (Sperber/Wilson 1995, p.58)

The assumptions can be more or less present to the speaker. In some cases they are completely present to the speaker as in the Ticket collector's answer in (33):

(33) *Passenger*: When does the train arrive at Oxford?

Ticket-collector: At 5:25.

Here the ticket-collector wants to make manifest the single assumption that the train will arrive at Oxford at 5:25. This assumption is clearly present in the ticket-collector's mind. It is, however, a mistake to think that all communication is like this. Sperber & Wilson give

another example where none of the assumptions in **I** are directly listed (Sperber/Wilson 1995, p. 59): Mary and Peter are in a house by the seaside and Mary opens the window and ostensibly sniffs the seaside air. Mary need not intend to communicate any particular assumption, but rather all the assumptions which became manifest to her when she opened the window. If one were to ask Mary what she wanted to convey, one of the best answers she could give would be to say that she wanted to share an 'impression' with Peter. An 'impression' would best be described as “a noticeable change in one's cognitive environment, a change resulting from relatively small alterations in the manifestness of many assumptions” (Sperber/Wilson 1995, p. 59).

In ordinary cases, what the speaker wants to convey is partly clear, as in (33), and partly vague as in Mary's sniffing. Some of the assumptions in **I** will be clearly present to the speaker, while others will be more vaguely present in her mind. The general point is that one should not think that in communication there is always a single thing 'meant' (Kenyon 2005, p. 131). Instead of focusing on the content of the utterance, one would be better off focusing on the wider 'point' of the utterance (Kenyon 2005, p. 141).

6.2.1 Successful communication

At this point it might be a good idea to define exactly what is meant by 'successful communication'. When is information successfully transmitted from the speaker to the hearer? It should be clear from the last section that successful communication does *not*, in general, rely on the hearer grasping the exact same proposition as the speaker expressed³³. In some cases, the proposition expressed is not even part of what the speaker wants to convey (as in (31), which expresses something false). What seems to matter is that the hearer gets the 'point' of the utterance. This is what communicative success is about, and I will draw a clearer picture of this in this section.

³³ There might be some cases where successful communication *does* rely on grasping the same proposition. RT does not need to disagree with this, as will be shown in the end of this section.

What RT means by successful communication is hinted at in the statement quoted in section 3.1 which I repeat here:

[C]ommunication 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 1995, pp. 192-3)

Daniel Wedgwood interprets this into a definition of successful communication:

[C]ommunication 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. (Wedgwood 2007, p. 11)

Wedgwood's definition is certainly in the spirit of what Sperber & Wilson say, but I will claim that it is not an entirely correct description of what communicative success amounts to in RT. I will show why I think Wedgwood's definition is imprecise and then give a definition of communicative success without these flaws.

The first problem with Wedgwood's definition is that it allows too many cases to count as successful communication. Say A and B are on a camping trip. They are about to make dinner, but John, who is also with them, is not in sight. A wants to know what to do about dinner and asks B where John is. B then answers:

(34) He went to the bank.

B is thinking of the sand bank close to the camping site, but A doesn't know about this sand bank and mistakenly thinks that John has gone back into the city to get some money³⁴. A then thinks that they will have to wait a long time to make dinner. This is clearly a case in which most people will describe A as misunderstanding B. However, after the exchange they *do* share some more assumptions than before: A now knows that B knows where John is; They both know that John went away for some time; etc. If we were to use Wedgwood's definition, we would have to describe this case as a case where communication is successful.

The problem with Wedgwood's definition is actually even more severe; It would count *any*

34 This example also clearly shows that not even disambiguating is a straight-forward matter.

communication as successful. Recognizing the speaker's *communicative intention* is necessary for something to be called communication, and when the hearer recognizes the speaker's communicative intention they immediately have one more assumption in common than before, namely the assumption that the speaker has an *informative intention*. It should be clear then, that Wedgwood's definition is too weak; According to it, communication is always successful.

The real problem with the definition is that it sets a clear cut-off point between successful and non-successful communication. It is not always easy to say whether communication is successful or not; Let us say that the sand bank in the example above is really far away. So far, in fact, that it takes just as much time to go there as it would take to go into the city. In some ways it is clear that A misunderstands B when thinking that John went to get some money, but the purpose of the question was to figure out what to do about dinner. A asked the question in order to find out how long John would be away. After hearing (34) A knows this, and can wait the appropriate time before preparing the food. In this altered example it seems that the communication was at least a little more successful than it was before.

I think this shows that a better way to think about successful communication would be to think that, just as with the notions of *resemblance* and *relevance*, success comes in degrees. Communication can be more or less successful. Sperber & Wilson put it this way:

Instead of treating an assumption as either communicated or not communicated, we have a set of assumptions which, as a result of communication, become manifest or more manifest to varying degrees. We might think of communication itself, then, as a matter of degree. (Sperber/Wilson 1995, p. 59).

A slight alteration of Wedgwood's definition is enough to correct the flaws in it:

Communicational success (comparative)

Communication is successful *to the extent that* through it the hearer shares the assumptions the speaker wanted to convey.

The set of assumptions they can come to share are the assumptions described in the speaker's *informative intention*. This means that the more of these assumptions they share

after the speaker has uttered something, the more successful the communication was. As long as the hearer is aware of the speaker's intention to communicate, we can say there is *some* success. It does not follow that this definition has the same flaw as in Wedgwood's definition by allowing too much to count as successful communication. Because Wedgwood's definition sets a clear cut-off point between successful and non-successful communication one would, with his definition, have to say that the communication was successful, period. In cases where the hearer totally misunderstands what the speaker tried to convey it would be misleading to say that the communication was successful. With my definition, one can say that the communication did not accomplish a lot, but at least the hearer knew that the speaker was trying to convey *something*.

One might ask if my definition allows for completely successful communication. It seems to be the case that if the hearer and speaker share *all* the assumptions in the speaker's set of assumption **I**, then we would have complete success. If this is the case, why not just say that at this point communication is successful, and say that anything less than this is non-successful communication? First of all, this would rule out many cases of communication as unsuccessful, even though they intuitively are successful (like the "Holland is flat" example mentioned above). Second, as already mentioned, RT claims that this kind of 'perfect communication' is seldom aimed for in normal circumstances. One would not want a criterion that rules out normal communication as successful.

Should we not at least give a criterion for when communication is successful enough? A general answer to this cannot be given. How many assumptions needed to be shared between the speaker and hearer will depend on the context and what the speaker/hearer is using the information for. Sometimes a loose understanding of an utterance will be sufficient, as when someone utters (30) in order to suggest that Holland as a good place to go bicycling. Other times, the speaker and hearer might need to understand each other perfectly, that is, share all the assumptions in the speaker's set of assumptions **I**. This is the case in (33), where the only assumption the ticket-collector wants to share is that the train leaves at 5:25. The hearer will have to grasp this exact proposition in order for this to count

as successful communication.

If RT claims that communication can be successful when the speaker and hearer only grasp similar propositions, does this mean that it is committed to the Similarity View? In the next section I will show that this is not the case and point to the important difference between RT and SV.

6.3 Resemblance does not lead to the Similarity View

As mentioned in section 3.2.1, according to C&L RT commits itself to the 'Similarity View'. Committing to this view means that one is committed to say things like: 'A said that p' means the same as 'A said something similar to p', and 'A said what B said' means 'A said something similar to what B said' (Cappelen/Lepore 2008, p. 1034). Having seen how 'resemblance' plays a role in RT, it is not difficult to see why someone would think that this implies a commitment to SV. Believing this would, however, mean that one has missed an important difference between RT and SV.

Wedgwood points to this mistake. He says that C&L fail to understand that “RT is not a semantic theory ... As such it simply does not deal in objective truth in the way that semantic theories tend to” (Wedgwood 2007, p. 5). It is important to keep in mind that RT is a psychological theory, concerned with what happens in the speaker and hearer during communication. It is generally *not* concerned with the truth-conditions of the propositions involved. Since the success of communication does not generally rely on the hearer grasping the same proposition as the one the speaker expressed, RT does not have to say anything about whether the propositions the speaker and hearer grasp have the same truth-conditions. RT is not claiming that “A said that p' means the same as 'A said something similar to p’”, but it *does* claim that sometimes communication can be successful even if the hearer only grasps something similar to what the speaker expressed. Note that these two claims are entirely distinct and that the latter does not imply the former. The hearer would,

in some objective sense, ascribe a false belief to the speaker, but as long as it is similar enough (shares the right logical and contextual implications in the context) it does not matter. Wedgwood says:

C&L's criticisms look at content-sharing from the point of view of an omniscient third party. In this objective sense, intended meaning and recovered meaning may indeed be only similar. RT, however, is concerned with the addressee's understanding of what must be the intended meaning of an utterance. (Wedgwood 2007, p. 10).

Admitting this, however, is not the same as believing that 'A said that p' means the same as 'A said something similar to p', it is just admitting that sometimes communication might not work flawlessly. In "Loose Talk" Sperber & Wilson(1985/1986) state their view quite clearly:

Our approach handles loose uses without abandoning truth-conditional semantics. If we are right, loose uses are non-literal uses in the sense described above: they are based on resemblance relations among representations, and involve interpretive rather than descriptive dimensions of language use. When a proposition or concept is loosely understood, it is not (or at least it need not be) that it is a vague concept or proposition; it is not that a guarantee of approximate truth is given to this proposition at all. Instead, certain of its logical and contextual implications are taken to be accompanied by regular guarantees of truth, whereas others are simply ignored. Thus the truth-conditional relation between propositions and the states of affairs they represent remains unaltered (Sperber/Wilson) 1985/1986, p. 164).

In other words: Saying that two propositions resemble each other is not the same as saying that they have the same truth value.

Because these considerations show that in some cases (but not all) RT *is* committed to (RNSC), it is important to see why this does not necessary mean that the consequences mentioned by C&L (in section 3.2) are imminent. In the next section I will revisit C&L's examples and show how, given RT, we can account for them.

6.4 Examples of successful communication

In section 3.2 I listed some examples given by C&L which were supposed to be cases where content sharing is necessary. One example involved a CIA task force trying to decide

whether Jane is a spy. C&L claim it is necessary that the people involved in this case understands this question in the same way to be able to find a correct answer. This seems to be correct. However, as is the case with a lot of C&L's examples, it gives a slightly misleading version of how we actually communicate.

If a CIA task force were to investigate whether Jane was a spy, one could assume that what they were supposed to investigate was better specified. The question they were supposed to investigate would probably involve a description of what exactly they suspected Jane might be doing. “Is she giving away classified military information to a specified alien government?” “In what ways is this information obtained?” Etc. It is improbable that the only order the agents working for the task force receive is: “Is Jane a spy?”. If this *were* the only description they had for what they were supposed to investigate, they probably would *not* understand this question in the same way.

C&L might agree with this, but answer that these considerations does not disconfirm what they are claiming. One might be able to specify the order, but the agents would still have to understand *that* order correctly. This is correct, and RT does not deny it. However, when the example is stated in this way it is much easier to see how the Relevance Theoretical procedure can provide the hearers with a correct interpretation of the order.

Something C&L fail to notice is that people are generally aware of how the context may affect the interpretation of an utterance. They even claim that our prima facie intuition is that most words are context-stable³⁵. I think this is wrong. My claim is that people *are* aware of how the change of context may influence their utterances. Take the following example: A is watching 12-year olds playing basketball, and B utters:

35 This is clear in the introduction of *Insensitive Semantics* (2005). They say that they sometimes give their students a list of indexicals, like 'I', 'You', 'now' and 'that', and then ask the students if they think 'penguin' is like that. Most students say “no”, and this is supposed to show that people in general think that words other than the indexicals are context-stable. Bezuidenhout (2006) correctly points out that this only shows that people do not think of 'penguin' as an indexical.

(35) Peter is really tall

because Peter is 180cm which is really tall for a 12-year old. A would not say that:

(36) B said that Peter is really tall

in a context where the height of adult basketball players is contextually salient, unless she is trying to mislead her audience. She might utter (36), but then she would add something like:

(37) but he was talking about 12-year old basketball players

Being aware of how utterances can mean different things in different circumstances is what makes people rephrase themselves if the context requires it. In general, the examples given by C&L would not be problematic in normal communication, because no one would utter things like “Naomi said that p” unless p was something that would mean the same thing in both contexts.

In addition, the awareness of how we might understand utterances differently sometimes makes us ask questions about other people's utterances. It is quite common to ask people to be more precise: “I heard that you said so-and-so, but does this mean that such-and-such is the case?” An agent wondering about what exactly he is supposed to find out to resolve the question of whether Jane is a spy can request for a more specified order. Normal communication rarely involves only one, short, utterance.

Furthermore, it is not true that grasping the same proposition is always necessary in order for the cases presented by C&L to work; Sometimes resemblance will do. In situations where the point of the utterance is not given directly by the proposition expressed, resemblance can be sufficient. This is shown in the following example: A says to a group of people:

(38) Bring me some money from the bank

A is thinking of the bank around the corner, but some of the people in the group do not know about that bank, and walk a couple of blocks in the opposite direction to the second-

to-nearest bank. All the people in the group addressed by A go to a bank, withdraw money and give them to A. As long as the point of the order was to get A some money, it does not really matter from which bank the money came. In this case some of the people in the group grasped only a similar proposition than what was expressed by A. Still, they were able to fulfill A's request.

All this said, misunderstandings *do* occur. Since there is no fail-safe procedure for interpreting someone's utterance, one can never be completely sure that the audience grasps the intended meaning of it. There is a possibility that someone in the CIA task force misunderstands whatever is meant by "Jane is a spy". He might even mistakenly arrest Jane because of this. Successful communication cannot never be *guaranteed*.

C&L are right in claiming that it is important to share content in a lot of situations. However we have no way of making absolutely sure that the right content is shared. All we can do is to try to provide enough information for the hearer to be able to make the right inferences. What the audience does with this information is out of our control.

7 Summary and conclusions

My aim in this thesis has been to defend RT from the claim that it cannot account for how people understand each other. I started in section 2 with a brief presentation of the general principles of RT. According to RT, a hearer always has to infer the speaker's meaning by analyzing the evidence given by the utterance, the context and other clues given by the speaker. Since communication is governed by a presumption of relevance, the hearer is justified in assuming that the first sufficiently relevant interpretation of the speaker's meaning he grasps is the intended one.

In section 3 I presented C&L's claim that RT can not explain how hearers grasp the speaker's meaning. They assert that the method provided by RT is too weak to account for this. Central to C&L's argument are the claims that RT is a version of what they call *The Similarity View*, and that RT is committed to the *Non-Shared Content Principle*. I then presented C&L's list of cases where, according to them, *content sharing* is necessary.

In section 4 I argued that C&L's Non-Shared Content Principle is imprecisely formulated. In its original form it obviously does not apply to RT. I then reformulated the principle to avoid the impreciseness in the original one, without weakening the force of C&L's argument. After that, I tried to show that RT is not overtly committed even to this, revised, principle.

In section 5 I uncovered an unarticulated argument in C&L's papers. This argument states that the only way a hearer can possibly grasp the proposition expressed, is if this proposition can be reached by semantics plus some minimal contextual influence. I then discussed the various uses of 'content' and proposed some new definitions in order to form a clearer view of the debate. I also argued that there is no level of content that has a guarantee of being shared between the speaker and hearer in communication. Communication is something that takes place at a risk, and one can never be completely sure that what the

speaker is trying to convey will get across. I also showed that C&L's own theory cannot do the work they believe is required from it. I concluded the section by saying that RT does not have to be committed to the view that the hearer only grasps a proposition similar to the one expressed by the speaker, but that RT claims that *sometimes* this *is* the case.

In section 6 I presented RT's view that *resemblance* is enough to explain how speakers and hearers understand each other. As long as the propositions they grasp are similar enough, this is often sufficient to be a case of successful communication. I also argued that one should not think of communication as speakers wanting to convey one specific thing, but that speakers are generally trying to convey a *point*. I then argued that, in spite of some similarities, the resemblance view is not a version of The Similarity View. In the lights of these considerations, I revisited the examples given by C&L that were supposed to be problematic for RT. I showed how RT can account for those examples without relying on the speakers and hearers grasping the same propositions.

My main argument in this thesis has been that C&L do not have sufficient reasons for their claim that RT cannot account for how a hearer is able to grasp the speaker's meaning. The relevance theoretical procedure *does* show how the hearer is being guided to the right proposition. However, one might argue that this does not show that there exist no other proposition that is similar enough to satisfy the constraints given in a certain context. RT does not *prove* that the hearer ends up with the exact same proposition as the speaker intended to convey. The only answer one can give to this is that, since there is no level of content that has a guarantee of being shared, there is *no way* to prove that the hearer grasps the exact same proposition as the speaker expressed. However, in many cases, grasping only similar propositions might be enough for the hearer to understand what the speaker meant.

8 Bibliography

- Bezuidenhout, Anne L. 1997. "The Communication of De Re Thoughts". In *Noûs* 31:2 (1997) pp. 197-225.
- Bezuidenhout, Anne L. 2006. "The coherence of contextualism: A reply to Cappelen & Lepore". In *Mind and Language* 2006.
- Borg, Emma. 2006. "Minimalism versus Contextualism in Semantics" (draft). URL: http://www.reading.ac.uk/AcaDepts/ld/Philos/borg/Minimalism%20vs%20Contextualism%20in%20Semantics_2.doc 08.05.2010.
- Cappelen, Herman & Lepore, Ernie. 2005. *Insensitive Semantics*. Blackwell Publishing.
- Cappelen, Herman & Lepore, Ernie. 2006. "Precis of *Insensitive Semantics*". URL: <http://folk.uio.no/hermanc/docs/PPRALL7.pdf> 08.05.2010
- Cappelen, Herman & Lepore, Ernie. 2007. "Relevance Theory and Shared Content" (draft). URL: <http://folk.uio.no/hermanc/docs/RelandSC4.pdf> 08.05.2010
- Cappelen, Herman & Lepore, Ernest. 2008. "Shared Content". In *The Oxford Handbook of Philosophy of Language*: 1020-1056. Oxford University Press. New York.
- Carston, Robyn. 2001. "Explicature and semantics". In S. Davis & B. Gillon (eds.) *Semantics: A Reader*. Oxford: Oxford University Press.
- Carston, Robyn. 2002. *Thoughts and Utterances*. Blackwell Publishing. UK.
- Carston, Robyn. 2007. "Linguistic communication and the semantics/pragmatics distinction". In *Synthese* (2008) 165. pp. 321-345.
- Carston, Robyn. 2009. "The Explicit/Implicit Distinction in Pragmatics and the Limits of Explicit Communication". In *International Review of Pragmatics* 1 (2009). pp. 35-62.

- Elugardo, Reinaldo. 2007. "Minimal Propositions, Cognitive Safety Mechanisms, and Psychological Reality". Draft. URL: <http://www.ou.edu/ouphil/faculty/elugardo/MinimalPropositions.pdf> 08.05.2010
- Hall, Alison. 2008. "Free Enrichment or Hidden Indexicals". In *Mind & Language*, 23. pp. 426-456
- Katz, J. 1981. *Language and other abstract objects*. Blackwell. Oxford.
- Kenyon, Tim. 2005. "Non-Sentences, Implicature, and Success in Communication". In *Ellipsis and Nonsentential Speech* (2005). pp. 131-147.
- Korta, Kepa & Perry, John. 2006. "Varieties of Minimalist Semantics". In *Philosophy and Phenomnological Research* Vol. LXXIII, No. 2, September 2006.
- Kriempardis, Eleni. 2009. "Shared Content as Speaker Meaning". In *Lodz Papers in Pragmatics* 5.2 (2009). pp. 161-190.
- Recanati, François. 2001. "What Is Said". *Synthese* 128. pp. 75-91.
- Recanati, François. 2004. *Literal Meaning*. Cambridge University Press. New York.
- Sperber, Dan & Wilson, Deirdre. 1985/1986. "Loose talk". In *Proceedings of the Aristotelian Society*, 86, pp. 153-171.
- Sperber, Dan & Wilson, Deirdre. 1995. *Relevance: Communication & Cognition (Second Edition)*. Blackwell Publishing. USA.
- Sperber, Dan & Wilson, Deirdre. 1998. "The Mapping Between the Mental and Public Lexicon". In *Thought and language*. CUP, Cambridge, 1998. pp. 184-200.
- Sperber, Dan & Wilson, Deirdre. 2002. "Pragmatics, Modularity and Mind-reading". In *Mind and Language* Vol. 17. pp. 3-23.
- Sperber, Dan & Wilson, Deirdre. 2005. "Pragmatics". In *The Oxford Handbook of Contemporary Philosophy*. Oxford University Press. USA. pp. 468-501.

- Sperber, Dan & Wilson, Deirdre. 2008. "A Deflationary Account of Metaphor". In *The Cambridge handbook of metaphor and thought*. Cambridge University Press. New York.
- Stanley, Jason. 2002. "Making it Articulated". In *Mind & Language*, 17. pp. 149-68.
- Wedgwood, Daniel. 2007. "Shared assumptions: semantic minimalism and Relevance Theory". URL: <http://www.ling.ed.ac.uk/~dan/CLRT-sg-web.pdf> 08.05.2010
- Wilson, Deirdre. 1995. "Is there a maxim of truthfulness?". In *UCL Working Papers in Linguistics 7*, pp. 197-212.
- Wilson, Deirdre & Sperber, Dan. 1993. "Linguistic Form and Relevance". In *Lingua* 90 (1993). pp. 1-25.
- Wilson, Deirdre & Sperber, Dan. 2002. "Truthfulness and Relevance". In *Mind, Vol 11*. pp. 583-632.
- Wilson, Deirdre & Sperber, Dan. 2004. "Relevance Theory". In *The Handbook of Pragmatics*. Blackwell Publishing. USA. pp. 607-632.